Banking on Artificial Intelligence: Information Hold-Up and Transformation of Relationship Lending*

Allen N. Berger [a]
University of South Carolina

Omrane Guedhami[b]
University of South Carolina

Zheyu Qi[c]
University of South Carolina

Raluca A Roman [d]
Federal Reserve Bank of Philadelphia

July 2026

Abstract

We find bank AI adoption raises corporate loan spreads by 9–13 bps, with propensity-score matching and Bartik-style instruments confirming robustness. AI converts soft information into proprietary hard information, deepening rather than reducing information asymmetry. Relationship price discounts erode at AI-adopting banks, yet relationships persist through non-price terms – smaller, shorter loans with fewer covenants – shifting function from spread reduction to rollover discipline. Machine learning (1995–2010) reprices public borrowers; deep learning and NLP (2011–2023) shift burdens toward private borrowers. Borrowers adopting AI face lower spreads, but this discount vanishes at AI-adopting banks, consistent with rent extraction over efficiency gains.

JEL Classification Codes: C45, G21, G24, D82, O31, O33.

Keywords: AI, artificial intelligence; bank innovation; syndicated loans; loan pricing; relationship lending; hard versus soft information; information rents; bargaining power; contract design.

* The views expressed are those of the authors and do not represent the views of the Federal Reserve Bank of Philadelphia, or the Federal Reserve System. [a] Berger: University of South Carolina, aberger@moore.sc.edu. [b]Guedhami: University of South Carolina, omrane.guedhami@moore.sc.edu. [c] Qi: University of South Carolina, zheyu.qi@grad.moore.sc.edu. [d]Roman: Federal Reserve Bank of Philadelphia, raluca.roman@phil.frb.org.

1. Introduction

Despite widespread interest in artificial intelligence (AI) transforming finance, limited evidence exists on how AI adoption by banks affects the fundamental economics of credit provision. Large U.S. banks have made substantial investments in AI-driven tools that are embedded in core lending activities. Many allocate billions of dollars annually to technologies able to process vast amounts of borrower information, enhance credit risk assessment, detect fraud, and strengthen regulatory compliance (Fares, Butt and Lee (2023); Crisanto, Prenio and Yong (2024); Gambacorta, Sabatini and Schiaffi (2025)). For borrowers, the consequences of bank AI adoption can be significant. Loan pricing, contract design, and the allocation of credit are all influenced by the extent to which lenders generate and use private information through these technologies. Yet despite the scale of these changes, there is little empirical evidence on how bank adoption of AI affects the cost of credit. We take up this question by studying how AI adoption shapes loan spreads and how it reshapes the traditional role of relationship lending based on soft information.

The theoretical prediction for how AI affects loan pricing is *ex ante* ambiguous. On the one hand, AI adoption may enable banks to generate proprietary information that is difficult for competing lenders to replicate, effectively deepening the information asymmetry between incumbent and outside lenders. This mechanism increases borrower switching costs and allows banks to extract greater informational rents (Rajan (1992); Petersen and Rajan (1995)). When a bank's AI system processes granular borrower data (including transaction patterns, supply chain linkages, and behavioral signals) it creates bank-specific knowledge that cannot be easily transferred or verified by other lenders. This *informational hold-up channel* predicts that AI adoption will increase the cost of credit, with banks capturing the technological rents. The hold-up logic is most transparent in established lending relationships, where the incumbent bank has accumulated a long history of proprietary borrower data that outside lenders cannot easily replicate. Yet even new credit relationships are susceptible: once AI begins processing a borrower's transaction data and behavioral signals, switching to another lender requires rebuilding that information stock from scratch, creating dynamic switching costs that reinforce the incumbent's bargaining advantage.

Alternatively, AI may reduce loan spreads through an *efficiency gains channel* by improving information quality, reducing screening costs, and enhancing monitoring capabilities

(Diamond, 1984, 1991). The bank may pass some of the cost savings to borrowers to help retain them as customers. Which force dominates – informational hold-up or efficiency gains – is ultimately an empirical question with significant implications for credit allocation, financial stability, and the welfare consequences of technological change in banking.

A second central question concerns how AI reshapes relationship lending. Do AI technologies make relationship lending stronger or weaker? Do banks change their business model away from relationship-based lending toward algorithmic-driven transactions? Banking relationships are central to credit market functioning, enabling banks to accumulate "soft" information about borrowers that is not easily verifiable or transferable, traditionally resulting in more favorable loan terms such as lower spreads (Berger and Udell (1995); Bharath, Dahiya, Saunders and Srinivasan (2011); Bolton, Freixas, Gambacorta and Mistrulli (2016)); Kysucky and Norden (2016); Beck, Degryse, De Haas, and van Horen (2018); Berger, Bouwman, Norden, Roman, Udell, and Wang (2024a,b, 2025). The emergence of AI creates two opposing forces that may fundamentally alter this equilibrium. On one hand, if AI converts soft, relationship-specific information into hard, verifiable data (for example, by codifying patterns that previously required human judgment) it may substitute for traditional relationship lending. In this substitution view, the tacit, judgment-based knowledge that gave relationship banks a competitive advantage over arm's-length lenders becomes algorithmic and therefore more easily replicated. Once soft information is encoded in hard data, the unique value of a long-standing relationship diminishes, and banks may shift toward transactional lending models that no longer require sustained borrower contact or proprietary information accumulation.

On the other hand, relationships may become more valuable under AI because long-term lending histories provide the rich, structured data that AI requires to achieve high predictive accuracy (Liberti and Petersen (2019)). In this complementarity view, the years of transaction records, payment histories, covenant compliance data, and qualitative assessments accumulated through a lending relationship are precisely the proprietary data assets that AI requires for high predictive accuracy. Banks with deeper relationships hold a structural advantage in AI deployment: their data are richer, longer, and unavailable to outside lenders, reinforcing rather than eroding the incumbency advantage. This complementarity predicts that relationships persist under AI, but the mechanism shifts from reducing adverse selection through soft-information screening to extracting rents through proprietary data advantages.

Empirically, distinguishing between substitution and complementarity is crucial for understanding how technological change affects the structure and stability of credit relationships and whether AI fundamentally alters banks' business models.

We test these competing hypotheses using a comprehensive empirical strategy that combines novel data on bank AI adoption with detailed loan-level information and addresses endogeneity through multiple identification approaches. We measure bank AI adoption using the USPTO Artificial Intelligence Patent Dataset (AIPD 2023 update), which classifies AI-related inventions in U.S. patents and pre-grant publications into eight technological categories (Giczy, Pairolero and Toole (2022)). We merge these data with syndicated loan information from the LPC DealScan database covering 1995 to 2023, bank characteristics from Bank Compustat, and borrower fundamentals from Compustat.

Our findings provide strong evidence for the informational hold-up hypothesis and document a fundamental transformation in how relationship lending operates. In our most saturated specifications with extensive controls and fixed effects, bank AI adoption raises loan spreads by approximately 12.6 basis points. For the average syndicated loan of $572 million, this translates to $720,000 in additional annual interest costs per loan. Instrumental variable estimates that address potential endogeneity are substantially larger, indicating spread increases of 80 to 108 basis points, suggesting our baseline estimates may be conservative. The effect is economically meaningful: aggregating our baseline estimates across all U.S. syndicated loans in our sample period, AI adoption redistributes approximately $1.5 to 2 billion annually from borrowers to banks, with substantially larger magnitudes implied by our IV estimates.

Turning to relationship lending, we find that AI does not eliminate relationships but alters how they function. The traditional price discount associated with relationship lending erodes at AI-adopting banks. The interaction between AI adoption and relationship indicators is negative (5 to 6 basis points), eliminating much of the traditional price benefit on loan spreads. However, relationships do not disappear. They migrate from price to non-price terms. After AI adoption, relationship borrowers receive significantly smaller loans (13% reduction), shorter maturities (8% reduction), and fewer covenants (4 percentage points lower). This reallocation suggests that AI reconfigures the function of relationships: rather than reducing information asymmetry and lowering spreads, relationships at AI-adopting banks facilitate tighter *ex ante* credit allocation and enhanced rollover discipline. AI converts soft, relationship-specific information into hard, verifiable data, but this conversion benefits the bank by enabling more precise control over credit supply rather than benefiting the borrower through lower prices. The evidence indicates that AI transforms the operational mechanism through which relationship lending generates value: from spread reduction toward enhanced credit control, a shift consistent with the informational hold-up interpretation and observable across our sample of syndicated corporate loans. .

The information hold-up mechanism generates sharp cross-sectional predictions that we confirm in the data. If AI creates bank-specific information advantages, effects should be stronger when: (i) AI processes previously unstructured information, expanding the set of borrowers subject to hold-up; (ii) borrowers have weak bargaining power and limited outside options; and (iii) banks have strong balance sheets and pricing power. We find precisely these patterns. During the first wave of AI adoption (1995–2010), machine learning tools excelled at processing structured, quantitative inputs – balance sheet ratios, income statement data, and credit metrics – precisely the type of financial information that publicly listed firms are required to disclose. AI gives the incumbent bank a new advantage not by accessing information that was previously unavailable, but by extracting non-obvious signals from shared public disclosures and encoding them in proprietary models that outside lenders cannot replicate.

As a result, Wave 1 AI primarily generates hold-up of public firm borrowers, whose structured disclosures are most amenable to machine learning. During the second wave (2011–2023), deep learning and natural language processing (NLP) dramatically expanded the set of information that AI can process: news text, earnings call transcripts, satellite imagery, supply chain linkages, and behavioral signals. This is the type of unstructured, idiosyncratic information that had traditionally been the exclusive domain of relationship lending with private firms, who lack the mandatory disclosure requirements of public companies. By codifying these previously opaque soft-information signals into proprietary hard data, second-wave AI deepens the bank's information advantage precisely where borrowers had historically been most protected by informational opacity. Consequently, Wave 2 AI shifts the hold-up burden dramatically toward private firm borrowers, who have no alternative disclosure channel through which to signal their quality to competing lenders and eliminates the residual pricing protection that established relationships had retained during the first wave, as the relationship interaction coefficient shifts by −22.5 basis points for private borrowers across waves.

Human-interaction AI technologies (natural language processing, computer vision, speech recognition) drive this second-wave effect, increasing private firm spreads by 20 basis points. Financially stronger, better-capitalized banks pass through larger spread increases (31 basis points vs. 15 basis points), while more profitable borrowers with stronger bargaining power face smaller increases (10 basis points vs. 15 basis points). These patterns are directionally consistent with AI amplifying informational advantages and enabling rent extraction, though formal difference tests are not uniformly significant, reflecting the inherent difficulty of detecting heterogeneity across subsamples

Several findings are not consistent with efficiency-based explanations. First, spreads rise rather than fall, inconsistent with efficiency gains being passed through to borrowers. Second, borrowers who themselves adopt AI (and thus generate more structured, transparent data that reduces screening costs) face 12 to 14 basis points lower spreads, but this discount disappears entirely at AI-adopting banks. If AI primarily improved screening efficiency and reduced information frictions symmetrically, both bank and borrower adoption should reduce spreads. Instead, we find stark asymmetry: borrower AI adoption reduces information costs and lowers spreads, while bank AI adoption creates proprietary information advantages and raises spreads. Third, financially stronger banks with greater pricing power exhibit larger spread increases, inconsistent with cost-based pricing but consistent with rent extraction. This asymmetry and the systematic heterogeneity patterns directly support the hold-up interpretation over efficiency-based alternatives.

Our identification strategy addresses the central concern that banks may adopt AI in response to changes in borrower quality or lending conditions. We employ three complementary approaches. First, our baseline specifications include borrower, lender, and year fixed effects along with extensive bank, borrower, and loan controls. Second, propensity score matching on observable bank characteristics (size, profitability, capital adequacy) confirms that even among observationally similar banks, AI adopters charge 28 basis points higher spreads. Third, we construct a Bartik-style instrumental variable that interacts banks' predetermined exposure to AI-intensive industries (measured over 1989 to 1994, before our sample period) with time-varying aggregate shocks to nonfinancial AI innovation. The instrument exploits the insight that banks with historical exposure to tech-sector borrowers respond more strongly to aggregate AI breakthroughs, but this interaction affects loan terms only through the bank's own adoption decision. First-stage *F*-statistics are very strong, and second-stage estimates indicate causal effects substantially larger than OLS estimates, suggesting that measurement error or omitted factors bias baseline estimates toward zero. Two complementary mechanisms explain this gap: classical measurement error in our patent-based AI measure attenuates OLS toward zero, and our instrument identifies a local average treatment effect (LATE) for complier banks for whom AI adoption generated the largest proprietary information advantages and thus the largest hold-up rents (Imbens and Angrist, 1994; Jiang, 2017)[1] Event study estimates using the Sun and Abraham (2021) interaction-weighted estimator document that the spread increases build gradually over the four years following adoption, consistent with proprietary information accumulating across successive lending cycles, before stabilizing, providing dynamic evidence that the hold-up mechanism deepens as the bank's information advantage grows.

These findings have important implications for credit markets and financial intermediation. The redistribution from borrowers to banks is not random: it falls disproportionately on private firms, less profitable borrowers, and those with weak bargaining power, precisely those most dependent on bank credit and least able to access alternative funding. The heterogeneity we document (stronger effects at high-market-power banks, larger impacts on inelastic borrowers, concentration of effects where information asymmetries are most severe, and the asymmetric treatment of bank versus borrower AI adoption) suggests that market power and hold-up mechanisms play significant roles. The patterns are more consistent with rent extraction than with symmetric efficiency improvements that benefit both banks and borrowers. Our results also speak to policy debates about AI adoption in finance. Banking regulators have expressed

[1] Under the LATE framework, our Bartik-style instrument identifies the causal effect for banks whose AI adoption was induced by their predetermined exposure to technology-intensive industries interacting with aggregate innovation shocks — banks that adopted AI in direct response to technological opportunity rather than as an endogenous response to borrower quality changes. These complier banks are plausibly those for whom proprietary information advantages are largest, which may explain why IV estimates exceed OLS. As Jiang (2017) documents, IV estimates in corporate finance average approximately nine times larger than corresponding OLS estimates, consistent with both LATE and attenuation bias. Our OLS and IV estimates therefore bracket the average treatment effect, with OLS a conservative lower bound

concern about AI creating opacity in risk assessment and potentially amplifying systemic risk. Our findings suggest a different concern: AI may be working exactly as designed (producing superior information that generates private returns) but the social returns may be smaller if gains manifest primarily as surplus redistribution rather than expanded access or improved allocation. The concentration of effects among private borrowers and relationship lending (settings where information asymmetries are most severe and borrowers have fewest outside options) suggests that AI may exacerbate rather than reduce existing frictions in credit markets.

This paper makes four contributions. First, we provide the first large-scale evidence on how incumbent banks' AI adoption affects the cost and structure of existing credit relationships. Recent work on fintech and machine learning has emphasized how new entrants using algorithmic credit models expand access and improve efficiency in competitive consumer credit markets (Berg, Burg, Gombovic and Puri (2020); Fuster, Goldsmith-Pinkham, Ramadorai and Walther (2022)). Our setting differs fundamentally: we study incumbent banks with established relationships, informational advantages, and market power in corporate credit markets characterized by high switching costs and relationship-specific investments. When AI is deployed in this environment, competitive pressure does not force pass-through of efficiency gains. Instead, AI deepens information asymmetries between incumbent and outside lenders, allowing banks to extract rents through higher spreads. The contrast is stark: Fuster, Goldsmith-Pinkham, Ramadorai and Walther (2022) find that machine learning adoption in competitive mortgage markets improves consumer welfare through better pricing and expanded access; we find that AI adoption by relationship lenders in concentrated syndicated corporate credit markets primarily redistributes surplus from borrowers to banks. This difference highlights that the welfare effects of financial technology depend critically on market structure, not just on the technology itself.

Second, we advance the literature on relationship lending (Petersen and Rajan (1994); Berger and Udell (1995); Boot (2000)) by documenting that technological change fundamentally restructures rather than eliminates relationship value. Relationships persist under AI but operate through different mechanisms: from reducing information asymmetry and lowering prices to facilitating monitoring and reallocating control rights through tighter, shorter contracts with fewer covenants. This transformation reconciles the persistence of relationship banking with the rise of hard-information technologies and shows that banks do not abandon relationship lending under AI but rather adapt how relationships function in their business models.

Third, we contribute to the literature on bank market power and loan pricing (Rajan (1992); Petersen and Rajan (1995); Santos and Winton (2008) ; Schenone (2010)) by documenting how technological innovation can amplify rather than reduce market power in credit markets. The concentration of spread increases at financially strong banks with high pricing power, and the directionally consistent exploitation of borrower-level variation in bargaining strength, suggest that AI adoption enables more effective rent extraction precisely where banks already possess market power.

Fourth, methodologically, we introduce a Bartik-style instrument for bank technology adoption that exploits predetermined variation in exposure to technology-intensive industries interacted with aggregate innovation shocks. This approach addresses endogeneity concerns that plague cross-sectional studies of bank technology adoption and may prove useful in other settings where financial institution behavior is potentially endogenous to market conditions.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on AI adoption in banking. Section 3 develops testable hypotheses. Section 4 describes data and methodology. Section 5 presents main results. Section 6 reports robustness tests. Section 7 concludes.

2. The Evolution and Role of AI in Modern Banking

Artificial intelligence (AI) is increasingly central to the transformation of modern banking. Bank of America recently announced a commitment of more than $4 billions to AI initiatives, focusing on task-specific tools for employees and delivering more personalized financial information to clients. In its commercial banking operations, the bank has already introduced an AI system that generates draft preparation materials for bankers ahead of client meetings.[2] In wealth management, the Merrill platform and the private bank unit employ AI-driven chatbots to enhance client service.[3] On the retail side, Bank of America's customer-facing assistant Erica has recorded over three billion interactions since its introduction in 2018, making it one of the most widely used AI

[2] Previously, bankers had to navigate various systems, including the client's portfolio, transactions, customer relationship management software as well as third-party research, a process which could take a junior banker many hours or even days.
[3] This tool helps advisers figure out how to proactively connect with clients about relevant investment opportunities.

tools in global banking. An internal employee-facing version, launched in 2020, is now used by more than 90 percent of the bank's 213,000 employees, underscoring the scale and integration of AI within the institution.[4]

As one of the largest and most representative bank holding companies (BHCs) in the U.S., Bank of America's extensive adoption and investment in AI demonstrates both the current breadth and the future potential of AI in the financial sector. Importantly, the role of AI extends well beyond internal efficiency gains. Across the industry, banks are increasingly applying AI to loan pricing, credit risk assessment, regulatory compliance, fraud detection, and customer engagement, highlighting the technology's wide-ranging impact on modern financial intermediation. In the aspect of credit risk assessment, banks are drawing on both traditional financial statements and alternative behavioral data to evaluate borrower risk with greater precision. At the same time, advances in machine learning have enabled large institutions to develop their own in-house risk assessment systems, capable of detecting suspicious transactions and potential fraud in real time[5]. Regulators are beginning to integrate AI into supervisory practices, using these tools to streamline compliance reporting and to monitor risks that may not be visible through conventional oversight mechanisms.[6]

The history of AI development in financial institutions can be traced back to the 1980s. The earliest applications in banking during the 1980s and early 1990s centered on symbolic AI and expert systems. These knowledge-based programs were used to flag suspicious transactions, support compliance, and experiment with the first forms of automated credit scoring. From the mid-1990s through 2010, banks increasingly adopted statistical machine learning models and big data approaches, which enhanced credit underwriting, enabled real-time fraud detection, and

---

[4] See https://newsroom.bankofamerica.com/content/newsroom/press-releases/2025/08/a-decade-of-ai-innovation--bofa-s-virtual-assistant-erica-surpas.html ; https://newsroom.bankofamerica.com/content/newsroom/press-releases/2025/04/ai-adoption-by-bofa-s-global-workforce-improves-productivity--cl.html ; https://www.americanbanker.com/news/how-bank-of-americas-erica-does-the-work-of-11-000-people

[5] JPMorgan Chase uses its proprietary AI platform (LOXM) for trade execution and also deploys AI-based models for risk management and compliance monitoring, part of its $18 billion annual tech spend. See https://www.businessinsider.com/jpmorgan-takes-ai-use-to-the-next-level-2017-8 ; https://www.mckinsey.com/industries/financial-services/our-insights/jpmorgan-chases-derek-waldron-on-building-an-ai-first-bank-culture

[6] Bank of England has suggested that AI-driven models may need to be incorporated into supervisory stress tests. See https://publications.parliament.uk/pa/cm5901/cmselect/cmtreasy/684/report.html ; https://www.reuters.com/sustainability/boards-policy-regulation/britain-needs-ai-stress-tests-financial-services-lawmakers-say-2026-01-20/

improved risk modeling. Since 2010, the focus has shifted toward deep learning, marked by rapid growth in the use of neural networks, natural language processing, and more recently, generative AI. These modern systems allow banks to process vast volumes of unstructured data, ranging from transaction histories and financial statements to emails and call transcripts, delivering predictive insights that go well beyond the capacity of traditional models.

With decades of development, AI has evolved from experimental applications into a foundational technology in banking. Each wave of innovation has reshaped how banks price loans, assess risk, monitor compliance, improve operating efficiency, and interact with customers, thereby transforming the broader financial market. Studying the incorporation of AI into credit markets provides a better understanding of how new technologies affect loan pricing, contracting, and the allocation of credit, as well as their broader implications for financial stability and market efficiency.

3. Development of the Empirical Hypotheses

In this section, we formulate our empirical hypotheses about the relation between bank AI adoption and subsequent effects on the borrowing cost on credit market and the relationship lending.

3.1 Effects on borrowing cost on credit market

Rajan (1992) shows that when lenders accumulate private information about their borrowers, they can capture informational rents by increasing switching costs and reducing borrowers' ability to obtain outside credit. The adoption of artificial intelligence (AI) amplifies this mechanism by enabling banks to process far richer and more proprietary information about borrower quality than traditional systems allow. AI-based credit technologies, ranging from machine learning and behavioral analytics to real-time transaction monitoring, generate unique, non-transferable insights into borrower risk and behavior. These informational advantages are difficult for competing lenders to replicate, effectively expanding the bank's informational monopoly (Fuster, Goldsmith-Pinkham, Ramadorai and Walther (2022); Petersen and Rajan (1994); Schenone (2010)) This process resembles the classic hold-up mechanism in relationship lending. Initially, AI improves efficiency by enhancing banks' ability to screen and monitor borrowers. Over time, however, the information produced by these systems becomes specific to the adopting bank and less transferable across institutions. Borrowers recognize that switching lenders would require rebuilding their data history elsewhere, which increases switching costs and weakens their bargaining position. Consequently, the informational rents created by these technologies accrue primarily to the incumbent lender.

In this setting, the benefits of technological progress are captured mainly by banks rather than borrowers. As emphasized by Stiglitz and Weiss (1981) information frictions affect not only which borrowers receive credit but also the price of that credit. By expanding their informational advantage, AI-adopting banks gain greater pricing power and can extract higher loan spreads. Based on this theoretical framework, we propose our first hypothesis:

Hypothesis 1a: Bank AI adoption will increase the cost of credit for borrowing firms.

While the preceding argument emphasizes that AI adoption may strengthen banks' informational advantage and increase borrowing costs through tighter credit standards and greater bargaining power, an alternative mechanism predicts the opposite effect. If the main consequence of AI adoption is to improve information quality rather than to amplify market power, then more accurate screening and monitoring should reduce uncertainty and the risk premiums embedded in loan pricing. In this view, AI functions as an information-enhancing technology that improves allocative efficiency in credit markets. Whether the overall effect of AI adoption raises or lowers borrowing costs therefore depends on which force dominates in practice, the potential for greater informational rents or the efficiency gains from more precise risk assessment.

From this perspective, AI adoption can reduce the cost of debt by improving information quality and lowering uncertainty. In classic models of financial intermediation, enhanced monitoring and reputation formation mitigate agency costs and lead to lower borrowing rates (Diamond (1984), (1991)). By processing large volumes of structured and unstructured data, AI strengthens banks' ability to observe borrower fundamentals directly, thereby reducing adverse selection and improving screening accuracy.

Recent empirical evidence supports this efficiency channel. Berg, Burg, Gombović and Puri (2020) show that digital footprints and alternative data enable lenders to predict default risk more accurately, expanding credit access for borrowers previously viewed as opaque. Fuster, Goldsmith-Pinkham, Ramadorai and Walther (2022) find that machine-learning credit models improve predictive power and pricing efficiency across credit markets, Together, these studies suggest that data-driven credit assessment can reduce informational rents and lower the cost of borrowing, particularly for firms that are more transparent. Consistent with an efficiency channel, recent work finds that bank AI adoption can improve ex-post loan outcomes. Ebrahimitorki and Kim (2025) show that AI adoption reduces banks' non-performing loan ratios, with larger gains in settings that require greater human interaction, suggesting complementarity between AI and human judgment.

Borrowers whose business models or technologies generate more standardized and transparent information, such as firms with extensive digital infrastructure or well-documented innovation activity, are especially likely to benefit from this mechanism. Greater transparency and verifiable data reduce lenders' uncertainty about borrower quality and therefore lower the compensation required for bearing credit risk (Boot and Thakor (2000), Liberti and Petersen (2019)). As a result, AI adoption can improve the precision of risk assessment and lead to more accurate pricing of credit, particularly for firms that are already information-rich and easier to evaluate. Based on this theoretical framework, we propose our second hypothesis

H1b: Bank AI adoption will decrease the cost of credit for borrowing firms.

3.2 Effects on bank AI adoption on relationship lending

Classic models of relationship lending emphasize that building long-term relationship allows banks to gather soft information which cannot be easily verified or transferred across institutions. Empirically, syndicated loan literature shows that borrowers with stronger relationships typically secure better loan terms, including lower spreads and fewer covenants (Ivashina and Scharfstein, (2010); Bolton, Freixas, Gambacorta and Mistrulli (2016) Liberti and Petersen (2019)). However, the rise of artificial intelligence may fundamentally alter this dynamic. Modern AI systems, particularly those built on machine learning and big data, generate increasingly rich and verifiable "hard" information by processing vast quantities of transaction records, digital footprints, supply-chain linkages, and behavioral patterns. In doing so, AI reduces banks' reliance on relationship-specific knowledge that traditionally supported loan pricing.

As AI is increasingly integrated in financial institutions, the informational advantages associated with relationship lending are likely to diminish. Borrowers who once benefited from long-standing relationships may no longer enjoy preferential loan terms once lenders can rely on algorithmic assessments. In this sense, AI adoption is transforming credit markets by substituting hard information for soft, thereby reducing the economic value of relationships in determining borrowing costs.

Hypothesis 2a: Bank AI adoptions substitutes the benefits of relationship.

An alternative mechanism suggests that artificial intelligence may complement the information produced through banking relationships. The effectiveness of AI models depends critically on the quantity, quality, and continuity of the data used for training and validation. Borrowers with long-standing relationships provide precisely such data that enhance the predictive power and reliability of machine learning algorithms Thus, rather than diminishing the value of relationships, technological adoption reinforces their role in providing reliable, relationship-based information for credit assessment. This view aligns with prior research showing that the integration of soft and hard information can enhance credit decision quality and that long-term relationships provide verifiable data that reduce informational asymmetry even as lending technologies evolve; (Stein (2002), Berger and Udell (1995); Bharath, Dahiya, Saunders and Srinivasan (2011)).

Therefore, AI adoption may complement the advantages of relationship lending and borrowers with rich lending histories can thus continue to enjoy favorable loan terms.

Hypothesis 2b: ***Bank AI adoption complements the benefits of relationship.***

4. Data and Sample Construction

Our data are drawn from multiple sources. To examine the effects of bank AI adoption on the cost of credit, we use the Loan Pricing Corporation (LPC) DealScan database covering the period 1995–2023. DealScan provides detailed information on syndicated loans to a wide range of corporate borrowers. We merge these loan records with bank-level data from Bank Compustat, borrower financials from Compustat, and bank AI adoption measures constructed from the USPTO Artificial Intelligence Patent Dataset (AIPD 2023 update; Giczy, Pairolero, and Toole, 2022).

The basic unit of observation is a loan facility reported in DealScan. Syndicated loans are large-scale credit contracts in which multiple lenders jointly provide financing to a single borrower. The average facility size in our sample is $572 million. While each facility has only one borrower, it may involve multiple lenders due to syndication. DealScan records the roles of participating lenders. We restrict the sample to lead lenders, since these banks play the central role in screening borrowers, setting contract terms, and monitoring loan performance (Bharath, Dahiya, Saunders, and Srinivasan, 2011). Following Ivashina and Scharfstein (2010), we define the lead bank as the "administrative agent" for a facility. If no lender is identified as the administrative agent, we classify as the lead bank the institution designated as the "agent," "arranger," "book-runner," "lead arranger," "lead bank," or "lead manager." For facilities with multiple lead banks, we retain the institution with the largest total assets.

Our measure of bank AI adoption is constructed from the USPTO Artificial Intelligence Patent Dataset (AIPD 2023 update). The AIPD identifies AI-related inventions in U.S. patents and pre-grant publications (PGPubs) using a machine learning–based patent landscaping methodology developed by the USPTO. The dataset covers 15.4 million USPTO patent documents published between 1976 and 2023, including 2.2 million additional documents released since January 2021. Within the AIPD, AI inventions are classified into eight component technologies: machine learning, computer vision, natural language processing, speech, evolutionary computation, AI hardware, knowledge processing, and planning and control.

To link bank AI adoption to syndicated loan activity, we first construct a master dataset of AI-related patents from the AIPD and merge it with PatEx application data to obtain key identifiers including application date, pre-grant publication number and date, patent number, and patent grant date. We then map patents to publicly traded firms through established crosswalks. Our primary link uses the KPSS patent–CRSP match (Kogan, Papanikolaou, Seru, and Stoffman, 2017), extended by authors up to 2023), which provides PERMNO identifiers that we convert to PERMCO and ultimately to Compustat GVKEY. Patent-level firm identifiers are then merged with the Compustat Bank dataset and the CRSP-FRB public banks list (updated to 2024) to flag banks.[7] Finally, we link these bank identifiers to lead lenders in the LPC DealScan database, following the definitions in (Ivashina and Scharfstein (2010) and Bharath, Dahiya, Saunders, and

[7] https://www.newyorkfed.org/research/banking_research/crsp-frb

Srinivasan (2011). This procedure yields a bank-level measure of AI adoption that can be directly matched to syndicated loan contracts.

Our main independent variable, *Bank AI*, equals one if the lead arranger has filed at least one AI patent within the five years preceding the loan syndication date, and zero otherwise.[8] Using a patent-based measure of AI capabilities is consistent with prior work that links AI-related innovation activity to real outcomes (Babina, Fedyk, and Hodson (2024)). To capture heterogeneity in AI technologies, we group the eight component technologies into three categories. *Foundation AI* includes machine learning, evolutionary computation, knowledge processing, and planning and control – technologies that form the core algorithmic infrastructure for AI applications. *Human-Interaction AI* includes computer vision, natural language processing, and speech – technologies designed to process unstructured data and human-generated content. *Hardware AI* consists of patents related to AI-specific hardware that enables computational scaling. Online Appendix Figure OA.1 plots the evolution of bank AI patenting over time, documenting that AI patent filings by banks remained near zero through the late 1990s, grew steadily through 2013, and surged dramatically after 2015, peaking above 1,600 patents annually around 2021, consistent with the technological waves that motivate our heterogeneity analysis.

For each loan in our sample, we obtain information on the lead bank, borrower, and loan characteristics from DealScan. Lead bank characteristics are merged from Bank Compustat data. Because both Bank Compustat and patent data cover publicly traded banks, our bank-level controls are available for the large bank holding companies (BHCs) that dominate the syndicated lending markets and account for the vast majority of syndicated loan volume during our sample period.

In some analyses, we use DealScan to distinguish between relationship and non-relationship borrowers. Following the literature, we classify a borrower as a relationship borrower if the lead bank has made at least one loan to the firm in the preceding five years, based on DealScan records of lead bank–borrower pairs.

Borrower financial information is drawn from Compustat, which provides accounting data for publicly traded U.S. firms. We match DealScan borrowers to Compustat firms using the

[8] Results are robust to using alternative AI adoption proxies calculated over three instead of five years preceding the loan syndication date.

GVKEY identifier and the Chava and Roberts, (2008) link file, updated through 2023. Consistent with prior work, we exclude financial firms (SIC 6000–6999) and non-U.S. firms (Bharath, Dahiya, Saunders, and Srinivasan, 2011). Borrower characteristics are measured as of the fiscal year ending immediately prior to the loan activation date.

A significant portion of DealScan borrowers are private firms for which Compustat data are unavailable. To preserve sample size and examine heterogeneous effects across public and private borrowers, we include private borrowers in our analysis with indicator variables for missing Compustat controls. This approach follows standard practice in the commercial lending literature and allows us to document economically important differences between public and private borrowers, recognizing that coefficient estimates for private borrowers rely on bank and loan characteristics rather than borrower fundamentals.

Loan-level characteristics, including loan spreads, maturity, collateral, covenant restrictions, and transaction fees, are obtained directly from DealScan for each facility during our sample window (1995–2023).

Table 1 presents variable definitions (Panel A) and descriptive statistics (Panel B). The average loan spread in our sample is 235 basis points over LIBOR, with substantial variation (standard deviation of 139 basis points). 80 percent of loans in our sample are originated by banks that have filed at least one AI patent in the preceding five years, reflecting the concentration of syndicated lending among large, technologically sophisticated institutions. 27 percent of loan facilities involve relationship borrowers. Among public borrowers with Compustat data, the average firm has total assets of approximately \$6.4 billion (log size of 1.86), modest leverage (8.5%), and positive profitability (ROA of 0.7%).

5. Empirical Results

5.1. Baseline Specification

To investigate the effects of AI adoption on banks' loan pricing, we estimate the following regression model (1):

$$\textit{Loan Spread Indicator}_{i,j,b,t} = \beta_0 + \beta_1 \textit{ Bank AI}_{b,t} + \beta_2 \textit{Bank Characteristics}_{b,t-1} + \beta_3 \textit{Borrower Characteristics}_{j,t-1} + \beta_4 \textit{Loan Characteristics}_{i} + \beta_5 \textit{Bank Fixed Effects}_{b} + \beta_6 \textit{Borrower Fixed}$$

$$Effects_j + \beta_7 Year\ Fixed\ Effects_t + + \beta_8 Loan\ Type\ Fixed\ Effects_i + \beta_9 Loan\ Purpose\ Fixed\ Effects_i + \mu_{i,j,b,t} \quad (1)$$

Our dependent variable is *Loan Spread Indicator* $_{I,j,b,t}$ is the cost of bank debt, *Loan Spread* is the all-in-drawn spread over LIBOR in basis points for loan *i* borrowed by firm *j* from lender *b* in year *t*. Our key independent variable, *Bank* $AI_{b,t}$, equals one if the lead arranger has filed at least one patent classified as artificial intelligence (AI) within the five years preceding the loan syndication date, and zero otherwise. We control for bank, borrower, and loan characteristics, as well as bank fixed effects (*Bank Fixed Effects* $_b$) borrower fixed effects (*Borrower Fixed Effects* $_j$) year fixed effects, loan type fixed effects, and loan purpose fixed effects to isolate the relationship between bank AI adoption and the cost of credit.

First, we control for bank characteristics (*Bank Characteristics* $_{b,t-1}$) that could influence the cost of bank debt. These variables include proxies for CAMELS exa mination ratings.[9] *Bank_Capital* is the ratio of bank equity capital divided by gross total assets (GTA)[10]; *Bank_Profitability* is return on assets, the ratio of the annualized net income to GTA); *Bank_Liquidity* is the ratio of bank liquid assets to GTA. In addition, we include *Bank_Size*, the natural logarithm of GTA. Banks in weaker financial condition may charge higher loan spreads to preserve profitability and meet short-term obligations. Large banks may be able to offer more favorable pricing since they benefit from economies of scale and pass on lower rates to borrowers. On the other hand, stronger and larger banks may also exercise greater market power, enabling them to charge higher loan prices.

Second, we control for borrower characteristics (*Borrower Characteristics* $_{j,t-1)}$ that could influence the cost of bank debt. These variables include *Borrower Size* (natural logarithm of the firm total assets (expressed in millions of dollars), *Borrower Book to Market (BTM,* the ratio of a firm's book value to its market value of equity, *Borrower ROA* (the ratio of net income divided by total assets), *Borrower Leverage* (the ratio of long-term debt plus debt in current liabilities over total assets), and *Borrower Tangibility* (the ratio of gross net property, plant, and equipment

[9] These proxies for CAMELS ratings are used in other studies, including Duchin and Sosyura (2014), Berger and Roman (2015), (2017), and Berger, Makaew, and Roman (2019).

[10] Gross total assets (GTA) equals total assets plus the allowance for loan and lease losses and the allocated transfer risk reserve (a reserve for certain foreign loans) as in Berger and Bouwman (2009).

(PP&E) scaled by total assets). For private borrowers lacking Compustat data, we include indicator variables for missing borrower characteristics (coefficients not reported for brevity).

Third, we control for loan characteristics *(Loan Characteristics $_i$)* that might be associated with the price of loans. We include *Loan Maturity*, the loan maturity in months, and *Log(Loan Size)*, the natural logarithm of the amount of a loan. Larger loans and longer maturities can capture economies of scale in bank lending, and may be inversely related to loan spreads.

Fourth, our fixed effects structure absorbs time-invariant heterogeneity at multiple levels. Bank fixed effects are dummies for each lead bank and capture unobserved heterogeneity across banks in lending practices, risk appetite, and competitive positioning. Borrower fixed effects are dummies for each borrower and capture unobserved heterogeneity across borrowers such as in credit quality, bargaining power, and bank relationships. Year fixed effects are dummies for the year of loan origination and capture unobserved temporal patterns in loan markets, including macroeconomic conditions and credit cycles. Finally, we include loan type and loan purpose fixed effects to absorb any differences in risk associated with different loan types and purposes, which may affect loans.[11] *Loan type fixed effects* include indicators for term loans, revolvers, and other loans. *Loan purpose fixed effects* include indicators for: acquisitions / takeovers, general purposes, recapitalizations, leveraged buyouts (LBOs), and other miscellaneous purposes (all other types of lending purposes). We also control for other loan contract terms, including indicators for whether the loan is secured by collateral and whether it contains financial covenants. Variable definitions for the analysis are reported in Table 1.

5.2. Main Results: AI Adoption and Loan Spreads

Table 2 presents baseline estimates of the effect of bank AI adoption on loan spreads, with *Bank AI* as the key variable of interest. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. Identification therefore comes from within borrower–lender variation over time, net of common time shocks and contract-type composition. Columns (1)–(3) use the full sample and sequentially expand the control set: lender characteristics (Column 1), borrower observables with missing-value indicators where unavailable (Column 2), and loan characteristics (Column 3). Columns (4)–(6) implement the

[11] The construction of the loan type and loan purpose dummies follows the methodology in Drucker and Puri (2009).

same specification sequence for public borrowers matched to Compustat. Columns (7)–(8) report corresponding specifications for private borrowers, beginning with lender controls and then adding loan characteristics.

Across all specifications, the estimated coefficients on *Bank AI* are positive and statistically significant. For the full sample and for public borrowers, coefficients are consistently significant at the 1% level. For private borrowers, the coefficients remain positive and significant at the 5% level. This pattern indicates that bank AI adoption is associated with higher loan spreads, with the effect most pronounced among public firms. In our preferred most complete specification, Column (3) for the full sample, Column (6) for public borrowers, and Column (8) for private borrowers, the estimated coefficients are 10.47 and 12.58 and 9.33 basis points, respectively, both significant at the 1% level.[12]

The economic magnitude is substantial. For the average syndicated loan facility of $572 million, a 12.58 basis point increase in spreads translates to approximately $720,000 in additional annual interest costs (0.001258×$572 million=$719,576). Aggregating across the U.S. syndicated loan market, which originates approximately $1.5 to $2 trillion annually during our sample period, AI adoption by banks redistributes an estimated $1.9 to $2.5 billion per year from borrowers to banks in our baseline specifications. Overall, the results are consistent with the informational hold-up hypothesis: banks extract rents by converting borrower information into proprietary knowledge that raises switching costs. In practice, these advantages translate into higher loan prices, as AI-adopting banks leverage improved screening, monitoring, and segmentation relative to outside lenders. This evidence aligns closely with H1a.

5.3. AI Adoption and Relationship Lending

Table 3 reports regression results on how bank AI adoption interacts with borrower–lender relationships in setting loan spreads. *Rel* is an indicator equal to one if the lead arranger has extended at least one loan to the borrower within the preceding five years. All regressions include borrower, lender, and year fixed effects, and loan type and loan purpose fixed effects. Columns

[12] The difference of 3.87 basis points between public and private borrowers is not statistically significant (*t*-stat=0.548), suggesting that AI adoption raises spreads for both borrower types, though with somewhat larger point estimates for public firms. These estimates likely represent lower bounds on the true causal effect: our instrumental variable analysis (Section 6.2) yields estimates that are much larger, suggesting that measurement error in AI adoption or omitted confounders attenuate rather than inflate our baseline estimates.

(1)–(3) present results for the full sample, public borrowers matched to Compustat, and private borrowers, respectively. All specifications include lender and loan controls. Columns (1) and (2) additionally include borrower controls.

Across specifications (1)–(3), the coefficient on *Bank AI* is positive and statistically significant, confirming that bank AI adoption is associated with higher borrowing costs. The coefficient on *Rel* is positive but statistically insignificant, suggesting that any average relationship-pricing effect is weak in this setting once the fixed effects and controls are absorbed. Most importantly, the interaction term *Bank AI* × *Rel* is negative and marginally significant at the 10 percent level in the full sample. The magnitude, roughly 5 to 6 basis points, is economically meaningful relative to the standalone AI effect and indicates that relationships partially offset, but do not eliminate, the AI-related spread premium.

To assess the total effect of AI adoption for relationship borrowers, we compute the net effect as the sum of the coefficients on *Bank AI* and *Bank AI* × *Rel*. For all borrowers, the net effect is 7.51 basis points (=12.75−5.23), which is statistically significant. For public borrowers, the net effect is 9.41 basis points (=14.81−5.40) and is statistically significant. For private borrowers, the net effect is 5.24 basis points (=11.22−5.98), which is not statistically distinguishable from zero. These results indicate that while AI adoption raises spreads even for relationship borrowers, the effect is attenuated relative to non-relationship borrowers, and the attenuation is strongest for private firms.

The differential treatment of relationship versus non-relationship borrowers is itself informative. The interaction term *Bank AI* × *Rel* is negative and similar in magnitude across all three samples (−5.23 to −5.98 basis points), indicating that relationships compress but do not eliminate the AI premium. The persistence of economically large net effects, particularly for public borrowers (9.41 bps), implies that even long-standing relationships provide limited insulation from AI-associated pricing.

Overall, the results suggest that AI adoption is reshaping relationship lending. Rather than being passed through to borrowers as lower prices, AI-enabled information appears to be internalized by adopting banks as proprietary advantage, increasing pricing power. This pattern is consistent with AI converting soft, relationship-specific information into hard, bank-owned signals that raise switching costs and facilitate rent extraction. This evidence is consistent with H2a.

5.4. Heterogeneity Analysis

5.4.1 AI Evolution Waves

AI has evolved through successive waves of technological innovation, each emphasizing different methods and applications. During the 1990s through 2010, banks adopted statistical machine learning models and big data methods that processed primarily structured financial information. Since 2010, the focus has shifted toward deep learning, with rapid advances in neural networks, natural language processing, and, more recently, generative AI. These technologies enable banks to extract information from unstructured sources such as textual filings, contracts, emails, and call transcripts.

We hypothesize that these technological waves have heterogeneous effects on different types of borrowers. First-wave machine learning tools processed structured financial data that public firms disclosed more comprehensively, potentially making public borrowers more susceptible to AI-driven information extraction. Second-wave deep learning and NLP tools process unstructured information that private firms, previously protected by opacity, now generate in electronic form. We therefore partition the sample into two periods: the first wave (1995–2010) and the second wave (2011–2023). The results are presented in Table 4.

Table 4 reports regression estimates of the effect of AI adoption across technological waves on loan spreads. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. Panels A and B report wave-specific estimates for the All, Public, and Private samples. Panel C reports tests of cross-sample and cross-wave differences in the *Bank AI* coefficient.

The results reveal a sharp reallocation of the AI spread premium across borrower types. In Panel A (first wave), AI adoption is associated with significantly higher spreads for public borrowers (15.2 basis points), while the corresponding estimate for private borrowers is small and statistically insignificant (2.0 basis points). In Panel B (second wave), the pattern reverses: the AI coefficient is economically large and statistically significant for private borrowers (19.4 basis points) but close to zero and statistically insignificant for public borrowers (0.6 basis points).

These findings are consistent with the view that the pricing effects of AI depend on which information margins the prevailing technology can most effectively exploit. In the first wave, when AI tools primarily leveraged structured inputs such as financial statements and market variables, public firms' richer and more standardized disclosures appear to have increased banks' ability to extract and monetize borrower information. In the second wave, advances in deep learning and natural language processing broadened information extraction to unstructured sources, increasing the scope for proprietary screening and monitoring of borrowers with limited standardized disclosure. Consistent with this shift, the spread premium associated with AI adoption concentrates among private borrowers in the post-2010 period. Panel C indicates that the public–private differential is economically meaningful in both waves.

The wave-specific pattern constitutes powerful evidence against alternative explanations. General improvements in bank efficiency or risk management would not produce a reversal in which borrower type bears the costs of AI adoption. Similarly, changes in borrower composition or loan characteristics absorbed by our fixed effects structure cannot explain why the public-private gap swings across technological regimes. The pattern is uniquely consistent with AI-driven information extraction: banks exploit whatever information their current technology can process, shifting from structured disclosures to unstructured private data as AI capabilities mature.

5.4.2 Technological Waves and Relationship Lending

The wave-specific patterns in Table 4 Panels A and B raise a natural follow-up question: does the role of relationship lending also evolve across AI waves? The hold-up mechanism predicts that as AI codifies soft, relationship-specific information into proprietary hard data, the traditional pricing protection that relationships provided should erode. Specifically, if Wave 1 machine learning tools primarily processed structured public disclosures, established relationships may have retained some residual insulating value during this period. In Wave 2, however, deep learning and NLP tools targeted precisely the unstructured, idiosyncratic information that relationship lending historically relied upon, suggesting relationship protection should diminish or disappear entirely in the second wave. Panel D tests this prediction by estimating AI adoption effects interacted with relationship indicators separately for each wave and borrower type.

The results are consistent with this prediction and reveal a striking evolution in the function of relationship lending across technological regimes. In Wave 1 (Columns 1–3), the interaction between bank AI adoption and relationship lending is negative for all firms (−5.1 basis points) and for public borrowers specifically (−5.2 basis points); however, neither interaction coefficient achieves conventional statistical significance on its own. The net effect for relationship borrowers [*Bank AI* + *Bank AI* × *Rel*] is positive and statistically significant for public borrowers in Wave 1 (9.4 basis points), indicating that relationship borrowers at AI-adopting banks still faced meaningfully higher spreads during this period, albeit partially moderated by the negative interaction. This pattern is consistent with established relationships partially offsetting but not eliminating the hold-up effect during Wave 1, when AI tools primarily extracted structured financial disclosures and the idiosyncratic soft information embedded in relationships retained some residual protective value.

In Wave 2 (Columns 4–6), this insulating role disappears entirely. The interaction term is statistically insignificant across all borrower types, and the point estimate for public firms reverses sign to +21.3 basis points, though estimated imprecisely. For private borrowers, where Wave 2 AI effects on the main Bank AI coefficient are largest (19.4 basis points), the relationship interaction is negligible; the net effect for private relationship borrowers in Wave 2 rises to 20.9 basis points (t=1.89), statistically indistinguishable from the main *Bank AI* effect and substantially larger than the corresponding estimate in Wave 1 (−1.6 basis points). The pattern is consistent with the view that second-wave deep learning and NLP technologies codified the soft, relationship-specific information that had previously insulated relationship borrowers, eliminating their pricing advantage. When AI can extract the same unstructured signals that relationship lending historically generated, the incremental value of a relationship in protecting against rent extraction approaches zero.

Panel E reports formal cross-sectional difference tests for both the interaction term and the net effect across waves and borrower types. For the interaction coefficient, the difference for public borrowers between Wave 1 and Wave 2 is −26.5 basis points (t=−1.54), marginally significant, confirming that the relationship interaction shifts meaningfully across waves for public firms from a small negative in Wave 1 to a large positive in Wave 2. For the net relationship borrower effect [*Bank AI* + *Bank AI* × *Rel*], the private firm cross-wave difference is −22.5 basis points (t =−1.75), significant at the 10% level: private relationship borrowers face a net spread increase 22.5 basis points larger in Wave 2 than in Wave 1, consistent with second-wave NLP and deep learning targeting precisely the unstructured information that had shielded private relationship borrowers during the machine learning era.

Together, the wave-relationship interactions in Panels D and E provide additional evidence for the informational hold-up mechanism and help explain why the erosion of relationship price discounts documented in Table 3 is concentrated in the modern AI era. The transition from Wave 1 to Wave 2 is not merely a change in which borrower types are most affected, it is a change in whether relationship lending itself retains any protective role in credit pricing.

5.4.3 AI Technology Categories

To examine which different AI technologies drive our results, we classify AI patents into three categories: *Foundation AI*, *Human-Interaction AI*, and *Hardware AI*. This classification reflects functional differences in how these technologies process information.

Foundation AI relies on core algorithms to process structured data. It includes patents related to: machine learning, evolutionary computation, knowledge processing, and planning and control. These technologies form the algorithmic core of AI by enabling banks to identify patterns in structured data and make predictions. These tools are particularly relevant for automating credit scoring and modeling borrower risk when firms provide standardized financial information.

Human-Interaction AI extracts insights from unstructured data. It includes patents related to: computer vision, natural language processing, and speech. These tools allow lenders to process unstructured data sources such as regulatory filings, loan contracts, earnings calls, and customer communications. This class of AI expands the informational basis of credit assessment and is especially important for evaluating private firms that often lack extensive disclosures.

Hardware AI refers to the computational infrastructure needed to implement AI effectively. It includes specialized chips, processors, and dataflow architectures. These technologies enable banks to scale AI applications with greater speed and efficiency.

Panel A of Table 5 reports the effect of different AI technology categories on loan spreads for the full sample period. *Foundation AI* adoption is associated with higher loan spreads for both public and private firms: 13.14 basis points for public firms and 9.89 basis points for private firms.

The magnitude is similar across borrower types, consistent with *Foundation AI*'s reliance on structured financial information that is available albeit in different forms for both public and private firms.

In contrast, *Human-Interaction AI* shows a strikingly different pattern. The effect is concentrated among private firms (7.55 basis points), with no significant effect on public borrowers (4.94 basis points, insignificant). This asymmetry is consistent with the information-extraction mechanism: *Human-Interaction AI* technologies which process unstructured text, speech, and images create informational advantages precisely where such unstructured data was previously difficult to analyze. Public firms already disclose extensive structured information; private firms do not, making them newly susceptible to NLP-based information extraction. For *Hardware AI* adoption, the pattern reverses: banks raise spreads more for private firms (14.77 basis points) than for public firms (8.81 basis points). The larger effect on private borrowers suggests that computational infrastructure investments yield greater informational returns where data was previously harder to process at scale.

Overall, Panel A reveals that all three AI technology categories raise loan spreads, but with heterogeneous incidence across borrower types. *Foundation AI* affects both similarly; *Human-Interaction AI* and *Hardware AI* affect private borrowers more strongly. These full-sample patterns motivate the wave-specific analysis that follows.

5.4.4 AI Technologies Across Waves

The full-sample results in Panel A may mask important temporal heterogeneity, given that AI technologies evolved substantially over our sample period. To examine how different AI technologies drive wave-specific patterns, we turn to Panels B through D of Table 5.

Panel B reports the impact of different AI technology categories on loan spreads during the first wave (1995–2010). For public firms, all three categories raise spreads significantly: *Foundation AI* increases spreads by 16.11 basis points, *Hardware AI* by 13.72 basis points, and *Human-Interaction AI* by 6.77 basis points. For private firms in the first wave, none of the AI technology categories has a statistically significant effect, with coefficients ranging from −2.11 to 7.15 basis points (all insignificant).

The cross-sectional difference tests at the bottom of Panel B formalize this pattern. The difference between public and private firms is statistically significant for *Foundation AI* (13.62 basis points), but not for *Human-Interaction AI* (8.88 basis points, insignificant) or *Hardware AI* (6.57 basis points, insignificant). The significant Foundation AI difference confirms that first-wave machine learning technologies designed for structured data primarily extracted information from public firms with rich financial disclosures.

Panel C reports results for the second wave (2011–2023). The pattern reverses dramatically. For private firms, all three AI categories now raise spreads significantly: *Foundation AI* by 19.42 basis points, *Human-Interaction AI* by 20.40 basis points, and *Hardware AI* by 22.90 basis points. For public firms in the second wave, none of the coefficients is statistically significant the point estimates range from 10.73 to 24.62 basis points, but standard errors are large, rendering all effects indistinguishable from zero.

The cross-sectional difference tests in Panel C confirm that public-private differences are not statistically significant in the second wave (*Foundation AI*: −8.69 bps; *Human-Interaction AI*: −3.10 bps; *Hardware AI*: 1.71 bps). This stands in stark contrast to the first wave, where *Foundation AI* showed a significant 13.62 basis point advantage for public firms. The reversal from significant public-firm effects in the first wave to significant private-firm effects in the second wave is the defining pattern of our wave analysis.

Panel D reports formal tests of cross-wave differences within each borrower type. For public borrowers, the changes across waves are uniformly insignificant: *Foundation AI* declines by 5.38 basis points (t=0.21), *Human-Interaction AI* increases by 10.53 basis points (t=0.69), and *Hardware AI* increases by 10.90 basis points (t=0.61). The lack of significant changes for public firms reflects the shift in AI's informational target rather than a decline in AI effectiveness.

For private borrowers, the cross-wave changes are uniformly positive and one is statistically significant. *Foundation AI* effects increase by 16.92 basis points (t=1.63, marginally insignificant), *Hardware AI* effects increase by 15.75 basis points (t=1.47, insignificant), but *Human-Interaction AI* effects increase by 22.50 basis points (t=2.44). This significant shift in Human-Interaction AI's impact on private borrowers is the sharpest statistical evidence that NLP and computer vision technologies which matured during the second wave enabled banks to extract information from previously opaque private firms.

Taken together, the evidence from Table 5 reveals that the economic consequences of AI adoption depend critically on the interaction between technological capabilities and borrower information environments. In the first wave, when AI tools primarily processed structured data, banks exploited the abundant structured information disclosed by public firms, leading to higher spreads for public borrowers. In the second wave, the emergence of *Human-Interaction AI* natural language processing, computer vision, and speech recognition enabled banks to process unstructured information from private firms. The "soft" qualitative information that previously protected private borrowers is now being converted into "hard" verifiable data, and the 22.50 basis point shift in *Human-Interaction AI*'s effect on private firms provides direct statistical confirmation of this mechanism.

5.4.5 Borrower AI Adoption

We now examine whether banks treat borrowers who have adopted AI adoption differently from non-adopters. As AI becomes an increasingly important general-purpose technology, firms that adopt AI may be perceived as more innovative and productive, with higher growth potential and stronger repayment capacity. Banks may also view technological alignment with their own AI adoption as a positive signal, leading to more favorable loan terms. (Chan, Lakonishok, and Sougiannis (2001); Eberhart, Maxwell, and Siddique (2004); Cornaggia, Mao, Tian, and Wolfe (2015); Cooper, Knott, and Yang (2022)).

To test this prediction, we interact *Borrower AI* with *Bank AI*. Table 6 Panel A reports the results. The analysis is restricted to public borrowers for whom patent data are available. The results show that borrowers who have adopted AI face significantly lower borrowing costs in the syndicated loan market, with spreads about 11 basis points lower (*Borrower AI*=−11.71). This effect is economically meaningful: AI-adopting borrowers pay approximately \$670,000 less in annual interest on the average loan facility, computed as (11.71 bps÷10,000)×\$572 million=\$669,812, consistent with borrower AI adoption generating transparency or quality signals that reduce information frictions.

However, the positive but insignificant coefficient on the interaction term suggests that banks do not offer additional pricing advantages to AI-adopting borrowers beyond this base effect. To assess whether AI-adopting borrowers face differential treatment at AI-adopting banks versus non-AI banks, we compute the implied effect of *Borrower AI* at AI-adopting banks, equal to *Borrower AI* + *Bank AI* × *Borrower AI*. At AI-adopting banks, the implied effect is −10.96 with t-test of 1.58. Thus, the Borrower AI discount remains economically similar at AI-adopting banks but is only weakly statistically distinguishable from zero based on the linear-combination Wald test.

This asymmetry is informative for interpreting our main results. *Borrower AI* adoption is associated with lower spreads, consistent with efficiency or transparency benefits that are at least partly passed through to borrowers. *Bank AI* adoption, by contrast, is associated with higher spreads in the baseline analysis, consistent with banks internalizing AI-enabled informational advantages as proprietary rents. The opposing signs help distinguish informational hold-up from a pure cost-efficiency channel.

5.4.6 Relationship Borrowers with and without AI adoption

An important question is whether a prior lending relationship affects this pattern. To test this hypothesis, we introduce two variables into the regression framework. *Borrower AI Rel* equals one if the borrower has filed at least one AI-related patent in the past five years and has an existing relationship with the lead lender, and zero otherwise. *Borrower Non-AI Rel* equals one if the borrower has not filed any AI-related patents in the past five years but has an existing relationship with the lead lender, and zero otherwise. Table 6, Panel B reports the results.

Table 6, Panel B shows that bank AI adoption is associated with higher spreads, consistent with the baseline findings. Relationship borrowers that adopt AI do not face an economically meaningful relationship premium: the coefficient on *Borrower AI Rel* is negative (about −4 basis points) and statistically insignificant. In contrast, relationship borrowers that do not adopt AI face significantly higher borrowing costs: the coefficient on *Borrower Non-AI Rel* is about +9.17 basis points and statistically significant.

Crucially, the wedge between these two relationship premium estimates is large and precisely estimated. The difference in relationship premiums, *Borrower AI Rel* − *Borrower Non-AI Rel*, is −13.02 in Column (2), with *t*-test of 3.00. This evidence indicates that borrower AI adoption materially alters the relationship-lending environment: AI-adopting relationship borrowers effectively avoid the relationship premium that applies to otherwise comparable non-adopting relationship borrowers. One interpretation is that borrower AI adoption improves outside options or bargaining position by making borrower quality more legible to competing lenders, thereby limiting the lead bank's ability to extract relationship rents.

5.4.7 Borrower Bargaining Power

More profitable firms and those with stronger balance sheets face better outside options and greater bargaining leverage with lenders. As a result, when banks adopt AI, borrowers with strong financial positions may be less affected by spread increases than financially weaker firms, which have fewer alternatives and weaker bargaining positions. To test this hypothesis, we stratify public borrowers using two measures of bargaining power. In Panel A, we classify borrowers by profitability (EBITDA/Assets): those below the sample median are designated *low bargaining power*, and those above the median are *high bargaining power*. In Panel B, we use leverage (debt/assets) as an alternative measure: borrowers with above-median leverage have weaker balance sheets and thus *low bargaining power*, while those with below-median leverage have *high bargaining power*. Table 7 reports the results.

Both panels reveal a consistent pattern: borrowers with weaker bargaining positions face larger spread increases following bank AI adoption. In Panel A, low bargaining-power borrowers (low profitability) experience significant spread increases of 15.3 basis points, compared to 9.8 basis points for high bargaining-power borrowers. The difference of 5.5 basis points implies that low-profitability borrowers face 56% larger spread increases than their more profitable counterparts. In Panel B, low bargaining-power borrowers (high leverage) face significant spread increases of 16.2 basis points, while high bargaining-power borrowers (low leverage) show a positive but statistically insignificant effect of 7.7 basis points. The difference of 8.5 basis points indicates that highly leveraged borrowers face more than double the spread increase of borrowers with stronger balance sheets.

The consistency across both measures strengthens the economic interpretation. Whether bargaining power is proxied by profitability or leverage, the pattern is the same: borrowers with weaker financial positions bear disproportionately larger costs when banks adopt AI. However, neither difference is statistically significant at conventional levels, reflecting the inherent difficulty of detecting heterogeneity in subsample comparisons. Detecting significant differences across subsamples requires substantially more statistical power than detecting significant main effects within subsamples. The fact that both coefficients are individually significant in Panel A, and the low-bargaining-power coefficient is highly significant in Panel B, confirms that AI adoption raises spreads across the distribution of borrower financial strength.

The economic interpretation is nonetheless clear. Banks appear to use AI-generated informational advantages to exert greater pricing power over borrowers with fewer outside options. Financially constrained borrowers, whether measured by low profitability or high leverage, are precisely those least able to credibly threaten to switch lenders, making them more susceptible to hold-up. The directional consistency across two distinct measures of bargaining power, despite the lack of statistical significance in formal difference tests, provides suggestive evidence that AI adoption amplifies the pricing consequences of borrower financial weakness.

5.4.8 Bank Financial Strength

While banks conduct extensive screening of borrowers during loan origination, borrowers likewise make deliberate choices about their lending partners. Firms often prefer banks with stronger balance sheets and lower risk profiles, viewing such institutions as more reliable counterparties that are less likely to face financial distress or curtail lending during downturns. This preference, however, can give financially stronger banks greater bargaining power in loan negotiations, allowing them to charge higher spreads. Conversely, weaker banks may need to offer more competitive pricing to attract and retain borrowers. If AI amplifies banks' ability to exploit informational advantages, the effect should be larger for banks with greater market power, that is, for financially stronger institutions.

To examine whether the effect of AI adoption on loan pricing varies with bank financial strength, we stratify by two measures in Table 8. In Panel A, we classify banks with below-median ROA as *financially weak* and banks with above-median ROA as *financially strong*. In Panel B, we stratify by capitalization, classifying banks with equity-to-asset ratios below the sample median as *financially weak* and those above the median as *financially strong*.

Panel A reveals an asymmetric pattern. Financially strong banks significantly increase loan spreads by 12.9 basis points following AI adoption, whereas financially weak banks show a positive but statistically insignificant effect (7.2 basis points). Only profitable banks appear able to convert AI adoption into higher pricing. The difference of 5.7 basis points (weak minus strong) is directionally consistent with stronger banks extracting larger rents but is not statistically significant. The lack of significance in the weak-bank coefficient rather than a significant difference between groups is the more informative finding: it suggests that financial strength may be a prerequisite for translating AI capabilities into pricing power.

Panel B examines capitalization and yields stronger results. Both well-capitalized and less-capitalized banks charge significantly higher spreads after adopting AI, but the effect is markedly larger among well-capitalized banks: 31.0 basis points compared to 14.7 basis points for less-capitalized banks. Well-capitalized banks charge more than double the AI premium of their less-capitalized counterparts. The difference of 16.3 basis points is economically substantial and approaches conventional significance levels (t =1.54).

The capitalization result merits interpretation. The pattern is inconsistent with cost-based pricing: if AI adoption represented a fixed investment with similar efficiency gains across banks, pass-through should be similar across bank types. Instead, the concentration of effects among well-capitalized banks suggests that financial strength enables more effective rent extraction, potentially through greater bargaining power or through the ability to absorb any borrower attrition resulting from higher prices. Well-capitalized banks may also have longer planning horizons, allowing them to invest in relationship-specific AI systems that deepen information asymmetries over time.

To summarize, both panels suggest that financially stronger banks extract larger AI-related premiums, though only the capitalization-based difference approaches conventional significance levels. The evidence is consistent with, though not definitive proof of, the hypothesis that bank market power amplifies the rent-extraction benefits of AI adoption.

5.4.9 Non-Price Contract Terms

If AI erodes the price discount associated with relationships, do relationships migrate to non-price contract terms? Table 9 examines this question. We analyze four contract dimensions: *Log(Loan Size)*, *Collateral*, *Covenant*, and *Log(Maturity)* as dependent variables.

Table 9 Panel A reports the direct effect of AI adoption on loan terms without relationship interactions. After a bank adopts AI, average loans are smaller, shorter, and less likely to carry financial covenants, with no detectable change in collateralization. For all borrowers, *Log(Loan Size)* falls by 3.4 percent, *Log(Maturity)* falls by 4.8 percent, and the probability of a covenant declines by 1.9 percentage points, while collateral shows no significant change (−0.6 percentage points, insignificant).[13]

The pattern differs across borrower types in economically meaningful ways. For public firms, loan size declines by 4.5 percent, maturity shortens by 4.0 percent, and covenants decline by 2.7 percentage points, with no change in collateral. For private firms, the size effect is smaller and insignificant (−1.9 percent), but maturity effects are larger (−5.6 percent) and covenant reductions are stronger (−3.0 percentage points). This heterogeneity suggests that AI adoption shifts contractual design toward tighter credit access, particularly through maturity for private firms, paired with simpler covenant structures for both borrower types.

Table 9 Panel B introduces relationship interactions to test whether AI fundamentally reshapes the non-price dimension of relationship lending. The *Bank AI* coefficient represents the effect for non-relationship borrowers; adding the interaction yields the effect for relationship borrowers. We report net effects for relationship borrowers with *t*-statistics from *F*-tests of joint significance.

For all borrowers, the results reveal a striking pattern. For non-relationship borrowers, AI adoption leaves facility size and collateral largely unchanged (coefficients of 0.017 and 0.001, both insignificant) but shortens maturity by 3.1 percent and reduces covenant incidence by 2.5 percentage points.

For relationship borrowers, the effects are dramatically larger and uniformly significant. Relationship borrowers at AI-adopting banks receive 12.4% smaller loans (net effect −0.132, t=−4.91), face maturities that are 8.0% shorter (net effect −0.083, t=−3.78), and encounter 4.1% fewer covenants (net effect −0.042, t=−3.12). Collateral requirements are also reduced by 2.2 percentage points, though this effect is only marginally significant (net effect −0.022, t=−1.73). The uniformly negative and statistically significant interaction terms across all four contract dimensions confirm that the shift from price to non-price mechanisms is markedly more pronounced for relationship borrowers than for transactional borrowers at the same AI-adopting banks.

[13] To calculate the percent change in the dependent variable, we use the following formula: $\%\Delta y=100*(\exp\beta -1)$.

The public-private comparison reveals important heterogeneity in how relationships transform under AI adoption. Among public firms, the interaction effects are large and statistically significant for size (−0.178), maturity (−0.080), and covenants (−0.035), with a marginally significant effect on collateral (−0.023). The net effects for public relationship borrowers are correspondingly large: loans are 15.0% smaller, maturities 8.6% shorter, and covenants 4.8 percentage points less likely, all statistically significant. Only collateral shows no significant net effect (−0.8 pp, t=−0.40).

Among private firms, the pattern differs notably. The interaction terms are significant only for size (−0.124) and collateral (−0.027), while the maturity and covenant interactions are essentially zero (−0.01 and 0.00, both insignificant). Nevertheless, net effects for private relationship borrowers remain significant: loans are 10.1% smaller (t=−2.94), maturities 6.2% shorter (t=−2.06), and covenants 3.0 percentage points less likely (t=−2.02). The private-firm net effects are driven primarily by the direct effect of AI adoption rather than by the incremental relationship interaction, suggesting that AI affects private borrowers more uniformly regardless of relationship status.

Relative to non-relationship borrowers, relationship borrowers at AI-adopting banks receive smaller and shorter loans, representing ex-ante tighter credit rationing, but they also face fewer covenants and reduced collateral requirements, which simplifies contracting. This pattern is consistent with a shift from price discounts toward dynamic rollover discipline. Banks substitute ongoing monitoring through rollover decisions for static contractual protections. The welfare implications for borrowers are *a priori* ambiguous and hinge on rollover certainty and future performance. Relationship borrowers gain contractual simplicity but face greater refinancing risk; the net welfare effect depends on whether banks use rollover discipline to support or to exploit their informational advantage**.**

To summarize, bank adoption of AI transforms soft-information relationship banking from a price-discount mechanism into a non-price control mechanism, compressing the marginal spread value of relationships while reallocating control to smaller, shorter, and simpler contracts governed by dynamic rollover discipline. The transformation is most complete for public firms, where all four non-price margins shift significantly for relationship borrowers. For private firms, size and collateral effects are pronounced, but maturity and covenant effects operate primarily through the direct AI channel rather than through relationship-specific interactions, consistent with the greater opacity of private borrowers limiting the precision with which banks can tailor non-price terms to relationship status.

6. Robustness Tests

6.1 Propensity Score Matching

One potential endogeneity concern is that facilities originated by AI-adopting (treated) banks are not randomly distributed across lenders, and observable lender fundamentals could drive both AI adoption and loan pricing. To mitigate this selection concern, we implement propensity score matching. Specifically, we estimate a logit model for the probability that a facility is originated by a treated lender using core lender characteristics: lender size, profitability (ROA), and capital adequacy. Using one-to-one nearest-neighbor matching with replacement, we construct a matched control sample of non-treated facilities that are comparable in propensity score to treated facilities.

Table 10 reports the matched-sample estimates of the effect of *Bank AI* on all-in-drawn spreads. The coefficient on *Bank AI* remains positive and statistically significant across matched specifications, with implied spread increases of 11.65 bps in the full sample, 10.63 bps for public borrowers, and 8.92 bps for private borrowers. Importantly, these magnitudes are economically close to the corresponding baseline estimates in Table 2, indicating that the loan-spread premium associated with bank AI adoption is not driven solely by differences in observable lender fundamentals that predict adoption. In combination with the saturated fixed-effect structure, the matched-sample evidence supports the interpretation that AI adoption is associated with systematically higher loan spreads even among facilities originated by observationally similar banks.

6.2 Instrumental Variables Estimation

A key concern in our setting is that bank adoption of AI may be endogenous to lending outcomes. For example, banks may innovate in AI in response to the composition of their borrower base or in anticipation of future lending conditions. To address this concern, we construct a Bartik-style instrument that exploits predetermined cross-sectional differences in bank exposures to AI-intensive borrowers and exogenous aggregate shocks in AI innovation outside of finance. Formally, our instrument for bank i in year t is defined as:

$$HistoricalTechExposureShock_{i,t} = Exposure_i \times Shock_{t-1} \tag{2}$$

where:

- $Exposure_i$: We measure each bank's ex-ante exposure to Al-intensive industries using the shares of loans of its syndicated loan commitments to SIC 7371 (Computer Programming Services), 7372 (Prepackaged Software), and 7373 (Computer Integrated Systems Design) during the 1989-1994 period. This window predates our estimation sample (1995-2023) and is held fixed throughout. These industries capture firms most directly associated with the development and integration of Al technologies. Using a fixed pre-period ensures that the exposure measure is predetermined and not influenced by bank' subsequent Al adoption decisions.
- $Shock_{t-1}$: The time-varying national shock is defined as the lagged growth in Al patents granted to nonfinancial firms (i.e., excluding banks and other financial institutions). Specifically, we compute:

$$Shock_{t-1} = \Delta log(NationalNonbankAIPatents_{t-1})$$

  This series captures broad technological waves in Al that originate outside of finance and diffuse into the banking sector. By construction, it excludes banks' own Al patenting activity and is lagged one year to mitigate simultaneity concerns.

- $Exposure_i \times Shock_{t-1}$ is the interaction of the two components described above and generates variation in AI adoption that is both cross-sectional (driven by banks' historical exposure to AI-intensive clients) and time-series (driven by exogenous national AI innovation). The identifying assumption is that banks with greater pre-period exposure to AI-intensive borrowers are more likely to adopt AI when nonfinancial AI innovation accelerates, but that this interaction affects lending outcomes only through its impact on bank AI adoption.

The empirical specification follows standard two-stage least squares (2SLS). The first stage regresses the indicator for bank AI adoption on the instrument $HistoricalTechExposureShock_{i,t}$ and bank-level controls, while the second stage instruments bank AI adoption in loan spread regressions estimated at the loan level.

To ensure that identification is not confounded by contemporaneous industry or macroeconomic shocks, we saturate the regressions with bank fixed effects, year fixed effects, and borrower × year fixed effects.[14] These controls absorb any time-invariant differences across banks, aggregate trends in AI development, and borrower-specific demand shocks. Consequently, the instrument identifies the differential sensitivity of banks with higher pre-period AI exposure to exogenous national AI innovation, consistent with best practices in Goldsmith-Pinkham, Sorkin, and Swift (2020).

We estimate the first stage at the bank–month level, consistent with the level at which the endogenous variable AI adoption is defined. Aggregating at this level avoids mechanical inflation of precision that would arise if the same bank–month observation was replicated across multiple loan-level entries. The second-stage regressions are then estimated at the loan level, linking bank-level adoption to facility-level loan terms.

Table 11 reports instrumental-variable estimates of the effect of bank AI adoption on loan pricing. Panel A presents the first-stage results, showing that *TechShock*, defined as the interaction of a bank's predetermined exposure to AI-intensive borrowers (measured over 1989–1994) and lagged growth in nonfinancial AI patenting, predicts subsequent bank AI adoption. The coefficients on both versions of the instrument are positive and statistically significant, indicating that banks with greater pre-period exposure to AI-related industries are more likely to adopt AI when external nonfinancial AI innovation accelerates. The *Kleibergen–Paap* statistics reported in Panel B imply strong identification and reject weak instruments.

Panel B reports the second-stage results, which use predicted AI adoption to estimate its effect on loan spreads. The 2SLS coefficients are economically large, ranging from approximately 80 to 116 basis points in the full and public-borrower samples and statistically significant at conventional levels. For private borrowers, the point estimates are also large (about 70 to 93 basis points), but statistical significance depends on instrument specification: the indicator-based estimate is imprecisely estimated (t=1.13), while the continuous-instrument estimate is significant

[14] We do not include bank fixed effects in the first stage because identification relies on cross-sectional differences in historical exposure across banks. Bank fixed effects would absorb exactly the variation the instrument is meant to exploit, leaving the first stage unidentified. Since the historical exposure variable is time-invariant by construction, a specification with year effects and bank-level controls is the appropriate structure for recovering the relevant first-stage relationship.

(t =2.64). Because *Bank AI* is an indicator variable, these coefficients should be interpreted as the effect of moving from non-adoption to adoption (for compliers identified by the instrument), rather than as a one–standard deviation change. Relative to Table 2, the IV magnitudes are substantially larger, consistent with baseline OLS estimates being attenuated.

Taken together, the IV results reinforce our main conclusion that bank AI adoption increases the cost of corporate credit, consistent with Hypothesis 1a. The larger 2SLS magnitudes are consistent with attenuation from measurement error in patent-based adoption proxies and with endogenous adoption dynamics that compress spreads in OLS (for example, banks adopting AI in response to competitive pressure or borrower mix changes that would otherwise reduce margins). By isolating variation in adoption induced by predetermined exposure interacted with aggregate innovation shocks, the IV strategy more directly captures exogenous technology-driven adoption, and the resulting estimates suggest that AI meaningfully strengthens lenders' informational advantage and pricing power.

As supplementary evidence on dynamics, Online Appendix Section OA.1 reports Sun and Abraham (2021) IW-weighted event study estimates around banks' first AI adoption year. Post-treatment coefficients rise monotonically from 29.8 basis points at $t = +1$ to 57.3 basis points at t=+4, consistent with gradual accumulation of proprietary borrower intelligence and progressive rent extraction. We discuss the pre-treatment pattern and the limitations of the design in the appendix.

6.3 Distinguishing AI from General Technological Sophistication

A potential concern with our identification strategy is that AI patenting may simply proxy for general bank technological sophistication rather than AI-specific information advantages. Banks that patent AI technologies may also be more innovative along other dimensions, and these broader capabilities rather than AI itself could drive our results.

To address this concern, Table 12 presents a falsification test that directly compares the effects of AI patents versus non-AI patents on loan spreads. We define non-AI patents as those with AI content scores below 10 percent according to the USPTO AIPD classification. If our results reflect general bank sophistication, both AI and non-AI patents should have similar effects on loan pricing. Alternatively, if AI-specific technologies drive our results, AI patents should have significantly larger effects.

The evidence supports AI-specific mechanisms. For public borrowers, AI patents significantly raise spreads by 11.37 basis points, while non-AI patents have a statistically insignificant effect of 3.43 basis points. The difference of 7.94 basis points is marginally significant, indicating that AI patents have economically and statistically larger effects than non-AI patents. This pattern is consistent with AI technologies generating proprietary information advantages that enable rent extraction, rather than general bank sophistication driving higher spreads.

For private borrowers, both AI patents (7.71 basis points) and non-AI patents (6.08 basis points) have statistically insignificant effects, and the difference between them is not distinguishable from zero. This attenuated pattern is consistent with our heterogeneity analysis in Table 4, which documents that private firm effects concentrate in the second technological wave (2011–2023). When we average across the full sample period including the first wave when AI had minimal effects on private firms the estimates are attenuated. The wave-specific pattern documented in Table 4 provides the relevant identification for private borrowers: effects emerge precisely when AI technologies enabled processing of unstructured information, a pattern that general bank sophistication cannot explain.

6.4 Additional Robustness Checks

Tables 13 present additional robustness tests. First, Panel A shows that results are robust to clustering standard errors at the bank level, addressing potential within-bank correlation of residuals. The results are robust: the coefficient on *Bank AI* remains positive and statistically significant for all borrowers as well as for both public and private borrowers. The slightly larger standard errors under bank-level clustering do not materially affect our inferences, confirming that the baseline results are not driven by understated standard errors.

Panel B addresses the concern that our results may be driven by a small number of large, technologically sophisticated banks that dominate the syndicated loan market. To address this possibility, we reestimate our baseline specification after excluding the three largest lenders by loan volume in our sample. The results strengthen rather than weaken when dominant lenders are excluded. For public borrowers, the coefficient on *Bank AI* increases from 12.58 to 13.73 basis points, and for all borrowers combined, the effect increases from 10.85 to 13.08 basis points. For private borrowers, the coefficient is 8.63 basis points, similar in magnitude to the baseline though less precisely estimated given the reduced sample size. These results confirm that our findings are not driven by idiosyncratic pricing behavior among the largest syndicated lenders.

Finally, Panel C examines whether the intensity of AI adoption matters, and reports estimates using the share of AI patents out of total patents filed by the bank over the preceding five years instead of an indicator variable for any AI patent in the preceding five years. The results confirm that greater AI intensity is associated with higher loan spreads. For public borrowers, a one-standard-deviation increase in AI patent share raises spreads by 7.88 basis points. For all borrowers, the effect is 6.09 basis points. For private borrowers, the coefficient is positive (4.95 basis points) but imprecisely estimated, again consistent with the concentration of private firm effects in the second wave documented in Table 4. These findings indicate that the relationship between AI adoption and loan pricing operates on both the extensive and intensive margins, with banks that invest more heavily in AI technologies extracting larger rents.

6.5 Summary of Identification and Other Robustness Tests

The robustness analysis yields four principal conclusions. First, effects for public borrowers are robust across all specifications: baseline OLS, propensity score matching, instrumental variables, falsification tests, bank-clustered standard errors, excluding dominant lenders, and AI intensity measures all yield statistically significant positive effects ranging from 8 to 116 basis points depending on identification approach.

Second, private borrower results are significant in baseline, PSM, and IV specifications but attenuated in some robustness tests. This pattern is fully consistent with our core heterogeneity finding: private firm effects concentrate in the second technological wave (2011–2023), so full-sample tests mechanically average over first-wave periods when AI had minimal effects on private firms. The wave-specific analysis in Table 4 provides the relevant identification for private borrowers—and there, effects are large (19.4 basis points) and highly significant.

Third, the falsification test directly addresses concerns that AI patenting proxies for general bank sophistication. For public firms, AI patents have significantly larger effects than non-AI patents (7.9 basis point difference), confirming that AI-specific technologies and not general innovation capacity drive our results.

Fourth, causal estimates from our Bartik-style IV approach are six to eight times larger than OLS estimates, with first-stage *F*-statistics ranging from 146 to 1,428, far exceeding conventional weak-instrument thresholds. This pattern suggests that measurement error or omitted variables bias baseline estimates toward zero, making our OLS estimates conservative lower bounds on the true causal effect.

Taken together, the identification evidence supports our baseline findings: bank AI adoption raises the cost of corporate credit through informational hold-up mechanisms, with effects concentrated among public firms in the first technological wave and private firms in the second wave.

7. Conclusion

Artificial intelligence (AI) adoption by banks is transforming the economics of corporate credit. While its stated goal often emphasizes improving efficiency in lending, enhancing risk management, and strengthening compliance, we find that AI adoption primarily imposes meaningful costs on borrowers in the credit market. Banks that adopt AI charge significantly higher loan spreads, raising borrowing costs by approximately 12.6 basis points in our baseline specifications and even larger in instrumental variable estimates. For the average syndicated loan of $572 million, this translates to $720,000 in additional annual interest costs per loan. Aggregating across the U.S. syndicated loan market, we estimate that AI adoption redistributes approximately $1.5 to 2 billion annually from borrowers to banks in our baseline specifications, with substantially larger magnitudes implied by our causal estimates. These effects are consistent with the view that AI technologies may allow banks to generate proprietary informational advantages and exercise greater bargaining power, deepening rather than reducing information asymmetry between incumbent and outside lenders.

The distributional consequences are notable as we show that our main effects are heterogeneous. Public firms bear the brunt of cost increases in the first wave of AI adoption (1995-2010), when statistical learning tools relied on structured financial information that publicly listed firms are required to disclose, while private firms are more affected in the second wave (2011-2023), when deep learning and natural language processing allowed banks to codify the unstructured, idiosyncratic soft information that had previously protected private borrowers through opacity, and exploit unstructured data. Different categories of AI tools also matter, with human-interaction AI particularly important for private borrowers. Financially stronger, better-capitalized banks tend to pass through larger cost increases, while borrowers with higher profitability and stronger bargaining power appear to face smaller increases. The costs of AI adoption may fall more on borrowers with weaker bargaining power precisely those most dependent on bank credit and least able to access alternative funding.

Perhaps most significantly, AI adoption substantially alters the role of relationship lending. Soft information is being converted into hard, verifiable data by evolved AI technologies. As a result, the traditional price discount associated with relationships erodes at AI-adopting banks, but relationships do not vanish. Instead, they migrate to non-price terms: loans become smaller, shorter, and less likely to include covenants, especially for relationship borrowers at AI-adopting banks. This transformation is consistent with AI converting soft, relational information into proprietary hard information that benefits banks through enhanced control rather than benefiting borrowers through lower prices. AI reshapes relationship lending from a price-discount mechanism into a control-rights mechanism emphasizing dynamic rollover discipline. Banks with stronger financial conditions pass through larger spread increases, consistent with AI amplifying rather than equalizing existing market power asymmetries between lenders and borrowers. To the extent that our results from syndicated corporate credit markets reflect broader lending dynamics, they suggest that the adoption of AI by incumbent relationship lenders may reinforce rather than reduce existing inequalities in credit access.

These findings may have important implications for borrowers, banks, and policymakers. For borrowers, the evidence highlights that AI adoption increases the cost of credit and alters the structure of contracts in ways that eliminate traditional relationship benefits in pricing while shifting control to banks through tighter, shorter contracts. For banks, the results show that AI strengthens their ability to extract rents and reallocate credit terms, particularly from borrowers with fewer outside options. The stark asymmetry we document that borrowers who adopt AI face lower spreads, but this discount vanishes entirely at AI-adopting banks reveals that the welfare effects of AI may depend on who deploys the technology and the market in which it operates.

From a policy perspective, the results raise critical questions about the distributional consequences of financial technology. Our findings contrast sharply with recent evidence on competitive fintech entry, which shows that algorithmic lending by new entrants expands access and improves consumer welfare (Berg, Burg, Gombović and Puri (2020); Fuster, Goldsmith-Pinkham, Ramadorai and Walther (2022)). We find that when incumbent banks with market power and established relationships deploy AI, competitive pressure does not force pass-through of efficiency gains and instead AI appears to enable more effective rent extraction. If stronger banks can use AI to extract greater rents, and if weaker borrowers face disproportionate cost increases, the long-run effects of AI adoption may reinforce rather than reduce existing inequalities in access to credit. The concentration of effects among private borrowers and relationship lending, where information asymmetries are most severe and borrowers have fewest outside options, suggests that AI may exacerbate rather than mitigate frictions in credit markets.

Whether the credit reallocation we document reflects efficiency gains or welfare-reducing transfers is an important question that our analysis cannot definitively answer. On the other hand, to the extent that AI reduces mispricing and improves monitoring, these technologies may still contribute to financial stability. However, our evidence that spreads rise rather than fall, that financially stronger banks extract larger rents, and that the traditional benefits of relationship lending disappear in pricing while migrating to control terms, suggests that market power and rent extraction may play significant roles. Future research may examine whether AI adoption affects real outcomes such as investment, employment, and firm growth, and whether regulatory intervention to preserve competition in AI-driven credit markets is warranted.

References

Babina, T., Fedyk, A., He, A. and Hodson, J., 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151, 103745.

Beck, T., Degryse, H., De Haas, R. and Van Horen, N., 2018. When arm's length is too far: Relationship banking over the credit cycle. *Journal of Financial Economics*, 127(1), 174-196.

Berg, T., Burg, V., Gombović, A., and Puri, M. (2020). On the Rise of FinTechs: Credit Scoring Using Digital Footprints. *Review of Financial Studies*, 33(7), 2845–2897.

Berger, A. N., and Udell, G. F. (1995). Relationship Lending and Lines of Credit in Small Firm Finance. *Journal of Business*, 68(3), 351.

Berger, A.N. and Bouwman, C.H., 2009. Bank Liquidity Creation. *Review of Financial Studies*, 22(9), 3779-3837.

Berger, A., Bouwman, C., Norden, L., Roman, R., Udell, G., Wang, T., (2024a). Piercing through Opacity Relationships and Credit Card Lending to Consumers and Small Businesses during Normal Times and the COVID-19 Crisis. *Journal of Political Economy* 132 (2), 484-551.

Berger, A.N., Bouwman, C., Norden, L., Roman, R.A., Udell, G.F. and Wang, T., (2024b). Relationship Lending and Monetary Policy Shocks: Evidence from the US. Working Paper.

Berger, A., Bouwman, C., Norden, L., Roman, R., Udell, G., Wang, T., (2025). Is a Friend in Need a Friend Indeed? How Relationship Borrowers Fare during the COVID-19 Crisis. *Journal of Financial Intermediation* 63, 101150.

Berger, A.N., Makaew, T. and Roman, R.A., 2019. Do Business Borrowers Benefit from Bank Bailouts?: The effects of TARP on loan contract terms. *Financial Management*, 48(2), 575-639.

Berger, A.N. and Roman, R.A., 2015. Did TARP banks get competitive advantages?. *Journal of Financial and Quantitative Analysis*, 50(6), pp.1199-1236.

Berger, A.N. and Roman, R.A., 2017. Did Saving Wall Street Really Save Main Street? The Real Effects of TARP on Local Economic Conditions. *Journal of Financial and Quantitative Analysis*, 52(5), pp.1827-1867.

Bharath, S.T., Dahiya, S., Saunders, A. and Srinivasan, A., (2011). Lending relationships and loan contract terms. *Review of Financial Studies*, 24(4), 1141-1203.

Bolton, P., Freixas, X., Gambacorta, L., and Mistrulli, P. E. (2016). Relationship and Transaction Lending in a Crisis. *Review of Financial Studies*, 29(10), 2643–2676.

Boot, A. W. A., and Thakor, A. V. (2000). Can Relationship Banking Survive Competition? *Journal of Finance*, 55(2), 679–713.

Callaway, B., and Sant'Anna, P.H.C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.

Chan, L. K. C., Lakonishok, J., and Sougiannis, T. (2001). The stock market valuation of research and development expenditures. *Journal of Finance*, 56(6), 2431–2456.

Chava, S., and Roberts, M. R. (2008). How does financing impact investment? The role of debt covenants. *Journal of Finance*, 63(5), 2085–2121.

Cooper, M., Knott, A. M., and Yang, W. (2022). RQ Innovative Efficiency and Firm Value. In *Journal of Financial and Quantitative Analysis*, 57(5), 1649-1694

Cornaggia, J., Mao, Y., Tian, X., and Wolfe, B. (2015). Does banking competition affect innovation? *Journal of Financial Economics*, 115(1), 189–209.

Crisanto, J.C., Leuterio, C.B., Prenio, J. and Yong, J., (2024). *Regulating AI in the Financial Sector: Recent developments and Main Challenges*. Bank for International Settlements, Financial Stability Institute.

de Chaisemartin, C., and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–2996.

Diamond, D. W. (1984). Financial Intermediation and Delegated Monitoring. *Review of Economic Studies*, 51(3), 393-414.

Diamond, D. W. (1991). Monitoring and Reputation: The Choice between Bank Loans and Directly Placed Debt. *Journal of Political Economy*, 99(4), 689–721.

Drucker, S. and Puri, M., 2009. On Loan Sales, Loan Contracting, and Lending Relationships. *Review of Financial Studies*, 22(7), 2835-2872.

Duchin, R. and Sosyura, D., 2014. Safer Ratios, Riskier Portfolios: Banks′ Response to Government Aid. *Journal of Financial Economics*, 113(1), 1-28.

Eberhart, A. C., Maxwell, W. F., and Siddique, A. R. (2004). An Examination of Long-Term Abnormal Stock Returns and Operating Performance Following R&D Increases. *Journal of Finance*, 59(2), 623–650.

Ebrahimitorki, M., Kim, H.H., (2025). Harnessing Artificial Intelligence: Impact of AI Adoption on Bank Loan Performance. Working paper, University of South Carolina.

Fares, O. H., Butt, I., and Lee, S. H. M. (2023). Utilization of artificial intelligence in banking sector: A systematic literature review. *Journal of Financial Services Marketing*, 28(4), 835–852.

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2022). Predictably Unequal? The Effects of Machine Learning on Credit Markets. *Journal of Finance*, 77(1), 5–47.

Gambacorta, L., Sabatini, F. and Schiaffi, S., (2025). Artificial intelligence and relationship lending. Bank of Italy Temi di Discussione (Working Paper) No, 1476.

Giczy, A. V., Pairolero, N. A., and Toole, A. A. (2022). Identifying artificial intelligence (AI) invention: A novel AI patent dataset. *Journal of Technology Transfer*, 47(2), 476–505.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8), 2586–2624.

Imbens, G. W., and Angrist, J. D. (1994). Identification and Estimation of Local Average Treatment Effects. *Econometrica*, 62(2), 467–475.Ivashina, V., and Scharfstein, D. (2010). Bank lending during the financial crisis of 2008. *Journal of Financial Economics*, 97(3), 319–338.

Jiang, W. (2017). Have Instrumental Variables Brought Us Closer to the Truth? *Review of Corporate Finance Studies*, 6(2), 127–140.Kysucky, V. and Norden, L., 2016. The benefits of relationship lending in a cross-country context: A meta-analysis. *Management Science*, 62(1), 90-110.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological Innovation, Resource Allocation, and Growth. *Quarterly Journal of Economics*, 132(2), 665–712.

Liberti, J. M., and Petersen, M. A. (2019). Information: Hard and Soft. *Review of Corporate Finance Studies*, 8(1), 1–41.

Petersen, M. A., and Rajan, R. G. (1994). The Benefits of Lending Relationships: Evidence from Small Business Data. *Journal of Finance*, 49(1), 3–37.

Petersen, M. A., and Rajan, R. G. (1995). The Effect of Credit Market Competition on Lending Relationships. *Quarterly Journal of Economics*, 110(2), 407–443.

Rajan, R. G. (1992). Insiders and Outsiders: The Choice between Informed and Arm's-Length Debt. *Journal of Finance*, 47(4), 1367–1400.

Santos, J. A. C., and Winton, A. (2008). Bank Loans, Bonds, and Information Monopolies across the Business Cycle. *Journal of Finance*, 63(3), 1315–1359.

Schenone, C. (2010). Lending Relationships and Information Rents: Do Banks Exploit Their Information Advantages? *Review of Financial Studies*, 23(3), 1149–1199.

Stein, J. C. (2002). Information Production and Capital Allocation: Decentralized versus Hierarchical Firms. *Journal of Finance*, 57(5), 1891–1921.

Stiglitz, J.E. and Weiss, A., (1981). Credit rationing in markets with imperfect information. *The American Economic Review*, 71(3), 393-410.

Sun, L., and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.

TABLE 1
Variable Definitions and Descriptive Statistics

This table presents variable definitions (Panel A) and descriptive statistics (Panel B). The sample consists of syndicated loans from LPC DealScan over 1995–2023. Bank AI adoption is measured using the USPTO Artificial Intelligence (AI) Patent Dataset (AIPD 2023 update). Bank characteristics are from Bank Compustat. Borrower characteristics are from Compustat and are available only for public borrowers. All continuous variables are winsorized at the 1st and 99th percentiles.

Panel A: Variable Definition

Variable Name	Description
Key Dependent Variables	
Loan Spread	Loan spread is calculated as all-in-Drawn spread over LIBOR in LPC DealScan. (For loans not using LIBOR, LPC transforms the spread into LIBOR equivalent by adding or deducting a difference updated periodically).
Key Independent Variables	
Bank AI	Indicator equal to one if a lead arranger has filed at least one patent classified as artificial intelligence (AI) within the five years preceding the loan syndication date, and zero otherwise.
Control Variables: Bank Characteristics	
Bank Size	The natural logarithm of the bank GTA
Bank Liquidity	Liquidity proxied by the ratio of bank liquid assets to GTA
Bank Capital	The ratio of stockholders' equity to total assets of bank.
Bank Profitability	The ratio of net income to total assets of bank.
Control Variables: Borrower Characteristics	
Borrower Size	The natural logarithm of the firm total assets (expressed in millions of dollars).
Borrower BTM	The ratio of a firm's book value to its market value of equity
Borrower ROA	Net income divided by total assets
Borrower Tangibility	Gross net property, plant, and equipment (PP&E) scaled by total assets
Borrower Leverage	The long-term debt plus debt in current liabilities over total assets.
Control Variables: Loan Characteristics	
Loan Maturity	The number of months to maturity on the loan.
Log(Loan Size)	Natural logarithm of the amount of the loan facility (in millions of dollars).
Collateral	Dummy equal to one for loan facilities secured by collateral, and zero otherwise.
Covenant	Dummy equal to one for loan facilities contains covenants, and zero otherwise.
Control Variables: Year, Bank, and Industry Fixed Effects	
Year Fixed Effects	Dummies for each loan origination year.
Bank Fixed Effects	Dummies for each of the lead banks.
Firm Fixed Effects	Dummies for each of the borrowers.
Control Variables: Loan Type Fixed Effects	
Term Loans	Indicator for term loans: Term Loan (Regular; A through H) and Delay Draw Term Loan.
Revolvers	Indicator for revolving lines of credit: Revolver/Line of Credit, 364-day Facility, or Limited Line.
Control Variables: Loan Purpose Fixed Effects	
Acquisition	Loan purpose indicator for Acquisition or Takeover.
General	Loan purpose indicator for General Corporate purposes, Capital Expenditure, or Working Capital.
LBO	Loan purpose indicator for purpose of LBO/MBO.
Recapitalization	Loan purpose indicator for Recapitalization or Debt Repayment.
Miscellaneous	Loan purpose indicator for the reasons: Project Finance, Trade Finance, Equipment Purchase, Stock Buyback, IPO Related Financing, Exit Financing, Spinoff, Real Estate, Telecom Buildout.
Instrument Variables	
TechShock (Continuous)	The interaction between each bank's predetermined exposure to AI-intensive industries (measured over 1989–1994) and lagged national growth in nonfinancial AI patents.
TechShock (Indicator)	Dummy variable that equals to 1 if TechShock >0, otherwise 0.

Other Variables Used in Additional Analyses	
Borrower AI	Indicator equal to one if the borrowing firm has filed at least one patent classified as artificial intelligence (AI) within the five years preceding the loan syndication date, and zero otherwise.
Rel	Indicator that equals one if the lead lender has provided a previous loan to the borrower in the last 5 years.
Foundation AI	Indicator equal to one if the bank has filed at least one AI patent in the past five years related to machine learning, evolutionary computation, knowledge representation, or planning and control, otherwise 0.
Human Interaction AI	Indicator equal to one if the bank has filed at least one AI patent in the past five years related to computer vision, natural language processing, or speech, otherwise 0.
Hardware AI	Indicator equal to one if the bank has filed at least one AI patent in the past five years related to AI hardware, otherwise 0.
Private Borrower	An indicator equal to one if the borrower is not publicly listed, and zero otherwise.
Borrower AI Rel	Indicator equal to one if the borrower has filed at least one AI-related patent in the past five years and has an existing relationship with the lead lender, and zero otherwise.
Borrower Non-AI Rel	Indicator equal to one if the borrower has not filed any AI-related patents in the past five years but has an existing relationship with the lead lender, and zero otherwise.

Panel B. Descriptive Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
Main Variables					
Loan Spread	62,835	235.2	139.23	22.50	750.0
Bank AI	62,835	0.800	0.400	0.000	1.000
Lender Characteristics					
Bank Size	62,835	13.39	1.242	9.302	15.136
Bank Liquidity	62,835	0.110	0.071	0.006	0.329
Bank Capital	62,835	0.090	0.018	0.030	0.134
Bank Profitability	62,835	0.011	0.006	-0.009	0.030
Borrower Characteristics					
Borrower Size	62,835	1.860	3.164	0.000	10.099
Borrower BTM	62,835	0.289	0.785	0.000	4.027
Borrower ROA	62,835	0.007	0.053	-0.304	0.177
Borrower Tangibility	62,835	0.087	0.189	0.000	0.841
Borrower Leverage	62,835	0.085	0.179	0.000	0.826
Loan Characteristics					
Loan Maturity	62,835	46.92	22.291	5.000	102.0
Log(Loan Size)	62,835	5.259	1.521	1.099	8.706
Collateral	62,835	0.094	0.293	0.000	1.000
Covenant	62,835	0.271	0.445	0.000	1.000
Term Loans	62,835	0.361	0.480	0.000	1.000
Revolvers	62,835	0.565	0.496	0.000	1.000
Acquisition	62,835	0.127	0.333	0.000	1.000
General	62,835	0.677	0.468	0.000	1.000
LBO	62,835	0.052	0.221	0.000	1.000
Recapitalization	62,835	0.021	0.142	0.000	1.000
Miscellaneous	62,835	0.077	0.267	0.000	1.000
Other Variables					
Rel	62,835	0.267	0.443	0.000	1.000
Foundation AI	62,835	0.798	0.401	0.000	1.000
Human Interaction AI	62,835	0.667	0.471	0.000	1.000
Hardware AI	62,835	0.749	0.434	0.000	1.000
Private Borrower	62,835	0.629	0.483	0.000	1.000
Borrower AI	6,943	0.631	0.483	0.000	1.000
Borrower AI Rel	6,943	0.025	0.431	0.000	1.000
Borrower Non-AI Rel	6,943	0.015	0.365	0.000	1.000
Instrument Variables					
TechShock (Continuous)	62,835	0.001	0.034	-1.622	0.986
TechShock (Indicator)	62,835	0.164	0.370	0.000	1.000

TABLE 2
Bank AI Adoption and the Cost of Corporate Credit

This table reports regression estimates of analyzing the effect of bank AI adoption on loan spreads. The dependent variable is *Loan Spread*, the all-in-drawn spread over LIBOR in basis points. *Bank AI* equals one if the lead arranger filed at least one AI-related patent in the five years preceding loan syndication. Columns (1)–(3) report results for the full sample (All borrowers); columns (4)–(6) report results for public borrowers matched to Compustat; and columns (7)–(8) report results for private borrowers. Within each sample block, the specifications are nested: columns (1) and (4) include bank controls only; columns (2) and (5) add borrower controls; and columns (3) and (6) include the full set of controls (bank, borrower, and loan controls). For private borrowers, Compustat data are unavailable; therefore columns (7)–(8) include bank controls and loan controls, but not borrower controls. Bank characteristics are bank size, liquidity, capital adequacy, and profitability. Borrower characteristics are borrower size, book-to-market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, and covenant. All regressions include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. Definitions of all variables are in Table 1. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sample	All	All	All	Public	Public	Public	Private	Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread
Key Variables:								
Bank AI	10.822***	10.145***	10.471***	13.549***	11.524***	12.578***	8.946**	9.329**
	(2.754)	(2.701)	(3.631)	(3.865)	(3.706)	(3.865)	(4.055)	(4.503)
Bank Controls								
Bank Size	-11.935***	-10.257***	-9.612***	-10.485***	-7.37**	-5.374	-12.464***	-13.681***
	(2.857)	(2.803)	(3.278)	(3.84)	(3.684)	(3.79)	(4.426)	(4.758)
Bank Capital	-34.726	-88.125	-76.969	31.327	-61.922	-41.079	-136.056	-136.168
	(65.986)	(64.718)	(1.162)	(92.902)	(89.08)	(-90.136)	(95.302)	(98.783)
Bank Liquidity	22.426	24.602*	22.031	15.967	19.941	17.904	31.233	29.709
	(14.299)	(14.021)	(1.549)	(21.073)	(20.203)	(20.344)	(19.627)	(20.042)
Bank Profitability	-246.525**	-284.817***	-308.073***	-130.643	-195.515	-188.685	-414.759**	-475.653***
	(111.611)	(109.454)	(2.687)	(153.867)	(147.554)	(152.421)	(164.676)	(175.361)
Borrower Controls								
Borrower Size		-9.605***	12.858		-11.144***	-9.975***	—	—
		(1.101)	(0.325)		(1.160)	(1.221)		
Borrower BTM		-5.413***	8.908**		-5.496***	-5.164***	—	—
		(1.177)	(2.239)		(1.196)	(1.213)		
Borrower Leverage		154.447***	49.051***		156.975***	159.503***	—	—
		(5.957)	(11.362)		(6.103)	(6.206)		
Borrower ROA		-160.611***	-22.470		-158.873***	-151.14***	—	—
		(11.786)	(1.215)		(12.005)	(12.241)		
Borrower Tangibility		24.334***	-19.233		26.967***	29.538***	—	—
		(7.788)	(0.555)		(7.922)	(8.131)		
Missing Borrower Controls Dummies		Yes	Yes				—	—
Loan Controls								

Log(Loan Size)			-4.605***			-3.757***		-5.956***
			(6.797)			(0.940)		(0.984)
Loan Maturity			0.080***			-0.095**		0.242***
			(2.989)			(0.038)		(0.037)
Collateral			13.076***			13.327***		10.938***
			(7.653)			(2.136)		(2.885)
Covenant			-6.622***			-6.823***		-3.691
			(4.856)			(1.722)		(2.265)
Observations	54,487	54,487	52,041	23,011	23,011	22,145	31,476	29,896
R-squared	0.735	0.746	0.678	0.692	0.718	0.722	0.761	0.763
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FEs: Borrower, Lender, Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

TABLE 3
Bank AI Adoption, Relationship Lending, and the Cost of Corporate Credit

This table reports regression estimates of the interaction between bank AI adoption and borrower–lender relationships on loan spreads. The dependent variable is *Loan Spread*, the all-in-drawn spread over LIBOR in basis points. *Bank AI* equals one if the lead arranger filed at least one AI-related patent within five years preceding loan syndication. *Rel* equals one if the lead arranger provided at least one loan to the borrower in the preceding five years. The net effect for relationship borrowers is computed as the sum of the coefficients on *Bank AI* and *Bank AI × Rel,* with the corresponding *t*-statistic. All specifications include bank controls and loan controls, and borrower controls in the All and Public samples. Bank characteristics are bank size, liquidity, capital adequacy, profitability. Borrower characteristics are borrower size, book to market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, covenant. All specifications include borrower, lender, year, loan type, and loan purpose fixed effects. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Sample	(1) All	(2) Public	(3) Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread
Bank AI	12.748***	14.805***	11.219**
	(3.138)	(4.179)	(4.748)
Rel	3.158	3.466	4.265
	(2.747)	(3.464)	(4.423)
Bank AI × Rel	-5.234*	-5.397	-5.981
	(2.954)	(3.786)	(4.669)
Net Effect Relationship Borrowers			
[Bank AI + Bank AI × Rel]	7.52**	9.41**	5.24
[*t*-stat]	[2.160]	[2.120]	[0.950]
Observations	52,041	22,145	29,896
R-squared	0.738	0.722	0.763
Bank, Loan Controls	Yes	Yes	Yes
Borrower Controls	Yes	Yes	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes

TABLE 4
AI Adoption across Technological Waves

This table reports regression estimates of the effect of bank AI adoption on loan spreads across two technological waves. The first wave (1995–2010) corresponds to statistical machine learning methods that primarily processed structured financial information. The second wave (2011–2023) corresponds to deep learning and natural language processing that expanded banks' ability to extract signals from unstructured information. The dependent variable is *Loan Spread*, the all-in-drawn spread over LIBOR in basis points. *Bank AI* equals one if the lead arranger filed at least one AI-related patent within five years preceding loan syndication. Panel A reports results for the first wave; Panel B reports results for the second wave. Panel C reports differences in the Bank AI coefficients across subsamples and waves, together with the associated *t*-statistics. Panel D replicates the specifications in Panels A and B with the addition of the relationship indicator and its interaction with Bank AI, estimating whether established borrower–lender relationships moderate the spread effects of AI adoption within each technological wave. Panel E reports differences in the Bank AI coefficients across subsamples and waves, together with the associated *t*-statistics. All specifications include bank controls and loan controls, and borrower controls in the All and Public samples. Bank characteristics are bank size, liquidity, capital adequacy, profitability. Borrower characteristics are borrower size, book to market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, covenant. All specifications include borrower, lender, year, loan type, and loan purpose fixed effects. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Panel A: FIRST WAVE (1995–2010)

Sample	(1) All	(2) Public	(3) Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread
Bank AI	9.984***	15.200***	2.013
	(3.232)	(3.979)	(5.458)
Observations	29,057	13,794	13,384
R-squared	0.762	0.769	0.780
Bank, Loan Controls	Yes	Yes	Yes
Borrower Controls	Yes	Yes	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes

Panel B: SECOND WAVE (2011–2023)

Sample	(1) All	(2) Public	(3) Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread
Bank AI	18.395**	0.607	19.42**
	(8.381)	(31.276)	(9.05)
Observations	20,568	5,434	14,379
R-squared	0.764	0.742	0.777
Bank, Loan Controls	Yes	Yes	Yes
Borrower Controls	Yes	Yes	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes

Panel C: FIRST WAVE vs. SECOND WAVE, Bank AI Tests of Cross-Sectional Differences

Bank AI	Difference	*t*-stat
FIRST WAVE: Public vs. Private	13.19**	[1.95]
SECOND WAVE: Public vs. Private	-18.81	[-0.58]
Public: FIRST WAVE vs. SECOND WAVE	14.59	[0.46]
Private: FIRST WAVE vs. SECOND WAVE	-17.41*	[-1.65]

PANEL D: WAVES AND RELATIONSHIP INTERACTIONS

WAVE	FIRST WAVE (1995–2010)			SECOND WAVE (2011–2023)		
	(1)	(2)	(3)	(4)	(5)	(6)
Sample	All	Public	Private	All	Public	Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread
Bank AI	11.623***	14.602***	3.727	18.257**	-3.412	19.378**
	(3.46)	(4.145)	(5.956)	(8.552)	(24.365)	(9.281)
Rel	-0.050	-0.313	.223	-3.886	-20.052	2.668
	(3.19)	(3.638)	(5.783)	(7.153)	(16.389)	(7.914)
Bank AI × Rel	-5.061	-5.221	-5.369	6.171	21.309	1.516
	(3.423)	(3.991)	(6.077)	(7.469)	(16.81)	(8.351)
Net Effect Relationship Borrowers						
[Bank AI + Bank AI × Rel]	6.562	9.380*	-1.642	24.427**	17.896	20.894*
[*t*-stat]	1.75	2.12	-0.250	2.43	0.680	1.890
Observations	29057	15334	14509	20568	6229	14379
R-squared	0.772	0.761	0.779	0.766	0.741	0.777
Bank, Loan Controls	Yes	Yes	Yes	Yes	Yes	Yes
Borrower Controls	Yes	Yes	—	Yes	Yes	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes	Yes	Yes	Yes

Panel E: FIRST WAVE vs. SECOND WAVE, [Bank AI × Rel]Tests of Cross-Sectional Differences

[Bank AI × Rel]	Difference	*t*-stat
FIRST WAVE: Public vs. Private	-0.148	[-0.02]
SECOND WAVE: Public vs. Private	19.793	[1.05]
Public: FIRST WAVE vs. SECOND WAVE	-26.530*	[-1.54]
Private: FIRST WAVE vs. SECOND WAVE	-6.885	[-0.67]

[Bank AI + Bank AI × Rel]	Difference	*t*-stat
FIRST WAVE: Public vs. Private	11.02	[1.39]
SECOND WAVE: Public vs. Private	-3.00	[-0.11]
Public: FIRST WAVE vs. SECOND WAVE	-8.52	[-0.32]
Private: FIRST WAVE vs. SECOND WAVE	-22.54*	[-1.75]

TABLE 5
AI Technology Categories

This table reports regression estimates of the effect of AI technology categories on loan spreads. The dependent variable is *Loan Spread*, the all-in-drawn spread over LIBOR, in basis points. *Foundation AI* captures patents in machine learning, evolutionary computation, knowledge processing, and planning and control, which represent core algorithmic methods primarily suited to structured data. *Human-Interaction AI* captures patents in computer vision, natural language processing, and speech, which facilitate the extraction of information from unstructured sources. *Hardware AI* captures patents related to AI-oriented computing infrastructure. Each category indicator equals one if the lead arranger filed at least one patent in that category within the five years preceding loan syndication. Panel A reports results for the full sample period. Panel B reports results for the first wave (1995–2010). Panel C reports results for the second wave (2010–2023). Panels B and C also report Public–Private differences in the estimated category coefficients. Panel D reports cross-wave differences within borrower type. All specifications include bank controls and loan controls, and borrower controls in the All and Public samples. Bank characteristics are bank size, liquidity, capital adequacy, profitability. Borrower characteristics are borrower size, book to market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, covenant. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

Panel A: FULL SAMPLE PERIOD

Sample	(1) All	(2) Public	(3) Private	(4) All	(5) Public	(6) Private	(7) All	(8) Public	(9) Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread
AI Technology:	**Foundation AI**	**Foundation AI**	**Foundation AI**	**Human Interaction AI**	**Human Interaction AI**	**Human Interaction AI**	**Hardware AI**	**Hardware AI**	**Hardware AI**
Bank AI	11.146***	13.143***	9.893**	5.324**	4.938	7.545**	11.49***	8.812**	14.774***
	(2.857)	(3.778)	(4.349)	(2.595)	(3.508)	(3.785)	(3.121)	(4.018)	(4.858)
Observations	52,041	22,145	29,896	52,041	22,145	29,896	52,041	22,145	29,896
R-squared	0.739	0.722	0.763	0.738	0.721	0.763	0.738	0.721	0.763
Bank, Loan Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Borrower Controls	Yes	Yes	—	Yes	Yes	—	Yes	Yes	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: FIRST WAVE (1995–2010)

Sample	(1) Public	(2) Public	(3) Public	(4) Private	(5) Private	(6) Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread
AI Technology:	**Foundation AI**	**Human Interaction AI**	**Hardware AI**	**Foundation AI**	**Human Interaction AI**	**Hardware AI**
Bank AI	16.111***	6.772*	13.717***	2.496	-2.105	7.152
	(3.879)	(3.725)	(4.164)	(5.135)	(4.678)	(6.268)
Observations	13,794	13,794	13,794	13,384	13,384	13,384
R-squared	0.769	0.768	0.768	0.780	0.780	0.780
Bank, Loan Controls	Yes	Yes	Yes	Yes	Yes	Yes
Borrower Controls	Yes	Yes	Yes	—	—	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes	Yes	Yes	Yes

Tests of Cross-Sectional Differences FIRST WAVE

	Difference	*t*-stat
FOUNDATION AI: Public vs. Private	13.62**	[2.12]
HUMAN INTERACTION AI: Public vs. Private	8.877	[1.48]
HARDWARE AI: Public vs. Private	6.565	[0.87]

Panel C: SECOND WAVE (2011–2023)

Sample	(1) Public	(2) Public	(3) Public	(4) Private	(5) Private	(6) Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread
AI Technology:	**Foundation AI**	**Human Interaction AI**	**Hardware AI**	**Foundation AI**	**Human Interaction AI**	**Hardware AI**
Bank AI	10.728	17.300	24.616	19.42**	20.395**	22.902***
	(24.768)	(14.777)	(17.413)	(9.05)	(7.956)	(8.669)
Observations	6,189	6,189	6,189	14,379	14,379	14,379
R-squared	0.760	0.759	0.760	0.780	0.780	0.780
Bank, Loan Controls	Yes	Yes	Yes	Yes	Yes	Yes
Borrower Controls	Yes	Yes	Yes	—	—	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes	Yes	Yes	Yes

Tests of Cross-Sectional Differences SECOND WAVE

Bank AI	Difference	*t*-stat
FOUNDATION AI: Public vs. Private	-8.692	[-0.33]
HUMAN INTERACTION AI: Public vs. Private	-3.095	[-0.18]
HARDWARE AI: Public vs. Private	1.714	[0.09]

Panel D: Additional Tests of Cross-Sectional Differences Across Waves

Bank AI	Difference	*t*-stat
Public: FIRST WAVE vs. SECOND WAVE, FOUNDATION AI	5.38	[0.21]
Public: FIRST WAVE vs. SECOND WAVE, HUMAN-INTERACTION AI	-10.53	[-0.69]
Public: FIRST WAVE vs. SECOND WAVE, HARDWARE AI	-10.90	[-0.61]
Private: FIRST WAVE vs. SECOND WAVE, FOUNDATION AI	-16.92	[-1.63]
Private: FIRST WAVE vs. SECOND WAVE, HUMAN-INTERACTION AI	-22.50**	[-2.44]
Private: FIRST WAVE vs. SECOND WAVE, HARDWARE AI	-15.75	[-1.47]

TABLE 6
Borrower AI Adoption and the Cost of Corporate Credit

This table reports regression estimates examining borrower AI adoption and loan spreads. The sample is restricted to public borrowers for whom patent data are available. The dependent variable is *Loan Spread*, the all-in-drawn spread over LIBOR in basis points. *Bank AI* equals one if the lead arranger filed at least one AI-related patent within five years preceding loan syndication. *Borrower AI* equals one if the borrowing firm filed at least one AI-related patent within five years preceding loan syndication. Panel A examines the interaction between bank and borrower AI adoption. Panel B examines the interaction between borrower AI adoption and lending relationships. *Borrower AI* × *Rel* equals one if the borrower has adopted AI and has an existing relationship. *Non-AI Borrower* × *Rel* equals one if the borrower has not adopted AI but has an existing relationship. All specifications include bank, loan and borrower controls. Bank characteristics are bank size, liquidity, capital adequacy, profitability. Borrower characteristics are borrower size, book to market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, covenant. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. *t*-statistics based on heteroskedasticity-robust standard errors are reported beneath the coefficient estimates. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

Panel A: BANK VS. BORROWER AI ADOPTION

	(1)
Sample	Public
Dependent Variable:	Loan Spread
Bank AI	7.049
	(7.576)
Borrower AI	-11.709*
	(6.461)
Bank AI × Borrower AI	3.921
	(6.708)
Implied Effects	
Borrower AI at non-AI Banks	-11.709*
Borrower AI at AI Banks	-10.96
[*t*-test]	[-1.58]
Observations	5,681
R-squared	0.752
Bank, Borrower Controls	Yes
Loan Controls	Yes
FEs: Borrower, Lender, Year	Yes
FEs: Loan Type, Loan Purpose	Yes

Panel B: BORROWER AI AND RELATIONSHIPS

	(1)
Sample	Public
Dependent Variable:	Loan Spread
Bank AI	13.323**
	(6.664)
Borrower AI Rel	-3.861
	(2.706)
Borrower Non-AI Rel	9.168***
	(3.325)
Interpretation:	
AI-Adopting Relationship Borrowers	No premium
Non-AI Relationship Borrowers	~+9.2 bps***
Diff REL premiums (Borr.AI - non-AI)	-13.02***
[*t*-test]	[3.00]
Observations	5,681
R-squared	0.749
Bank Borrower Controls	Yes
Loan Controls	Yes
Borrower, Lender, Year FE	Yes
Loan Type, Loan Purpose FE	Yes

TABLE 7
Borrower Bargaining Power and the Cost of Corporate Credit

This table reports regression estimates of the effect of bank AI adoption on loan spreads, stratified by borrower bargaining power. The sample is restricted to public borrowers. Panel A stratifies by borrower profitability (EBITDA/Assets): *Low Bargaining Power* includes borrowers with profitability below the sample median, and *High Bargaining Power* includes borrowers with profitability above the median. Panel B stratifies by borrower leverage (debt/assets). The dependent variable is *Loan Spread*, the all-in-drawn spread over LIBOR in basis points. For each panel, we also report the difference in coefficients of the effects for low and high bargaining power borrowers. *Bank AI* equals one if the lead arranger filed at least one AI-related patent within the five years preceding loan syndication. The *LOW–HIGH* difference is reported at the bottom of the table, along with the associated *t*-statistic. All specifications include bank, loan and borrower controls. Bank controls are bank size, liquidity, capital adequacy, profitability. Borrower controls are borrower size, book to market, leverage, ROA, and tangibility. Loan controls are loan size, maturity, collateral, and covenant. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

PANEL A: BORROWER PROFITABILITY

Sample	(1) Public	(2) Public
Dependent Variable:	Loan Spread	Loan Spread
Bargaining Power:	**LOW**	**HIGH**
Profitability:	Below Median	Above Median
Bank AI	15.345**	9.816**
	(7.640)	(4.958)
Observations	7,813	13,592
R-squared	0.744	0.760
Bank Borrower Loan Controls	Yes	Yes
Borrower, Lender, Year FE	Yes	Yes
Loan Type, Loan Purpose FE	Yes	Yes
Differences LOW-HIGH	5.53	
[*t*-stat]	[0.61]	

PANEL B: BORROWER LEVERAGE

Sample	(1) Public	(2) Public
Dependent Variable:	Loan Spread	Loan Spread
Bargaining Power:	**LOW**	**HIGH**
Leverage:	Above Median	Below Median
Bank AI	16.15***	7.658
	(5.370)	(5.892)
Observations	14,510	7,001
R-squared	0.650	0.711
Bank Borrower Loan Controls	Yes	Yes
Borrower, Lender, Year FE	Yes	Yes
Loan Type, Loan Purpose FE	Yes	Yes
Difference WEAK-STRONG	8.492	
[*t*-stat]	[1.07]	

TABLE 8
Bank Financial Strength and the Cost of Corporate Credit

This table reports regression estimates of the effect of bank AI adoption on loan spreads, stratified by bank financial strength. The sample is restricted to public borrowers. Panel A stratifies by bank profitability (ROA): *Financially Weak* banks have profitability below the sample median; *Financially Strong* banks have profitability above the median. Panel B stratifies by bank capitalization (equity/assets). The dependent variable is *Loan Spread*, the all-in-drawn spread over LIBOR in basis points. For each panel, we also report the difference in coefficients of the effects for financially weak and strong banks. *Bank AI* is a dummy variable equal to one if a lead arranger has filed at least one patent classified as artificial intelligence (AI) within the five years preceding the loan syndication date, and zero otherwise. For each panel, the WEAK minus STRONG difference in the Bank AI coefficient is reported at the bottom of the table, along with the associated *t*-statistic. All specifications include bank, loan and borrower controls. Bank controls are bank size, liquidity, capital adequacy, profitability. Borrower controls are borrower size, book to market, leverage, ROA, and tangibility. Loan controls are loan size, maturity, collateral, and covenant. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. Variable definitions are provided in Table 1. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

PANEL A: BANK PROFITABILITY

Sample	(1) Public	(2) Public
Dependent Variable:	Loan Spread	Loan Spread
Financial Strength:	**WEAK**	**STRONG**
Profitability:	Below Median	Above Median
Bank AI	7.186	12.928***
	(12.293)	(4.816)
Observations	9,957	11,196
R-squared	0.763	0.747
Bank Borrower Loan Controls	Yes	Yes
Borrower, Lender, Year FE	Yes	Yes
Loan Type, Loan Purpose FE	Yes	Yes
Difference WEAK-STRONG	-5.74	
[*t*-stat]	[-0.43]	

PANEL B: BANK CAPITALIZATION

Sample	(1) Public	(2) Public
Dependent Variable:	Loan Spread	Loan Spread
Financial Strength:	**WEAK**	**STRONG**
Capital Ratio:	Below Median	Above Median
Bank AI	14.72***	30.973***
	(5.059)	(9.246)
Observations	11,647	9,663
R-squared	0.779	0.727
Bank Borrower Loan Controls	Yes	Yes
Borrower, Lender, Year FE	Yes	Yes
Loan Type, Loan Purpose FE	Yes	Yes
Difference WEAK-STRONG	-16.25	
[*t*-stat]	[-1.54]	

TABLE 9
Bank AI, Relationships, and Non-Price Loan Contract Terms

This table reports OLS estimates of the effect of bank AI adoption and relationship lending on non-price loan contract terms. The dependent variables are: *Log(Loan Size),* the natural logarithm of the loan facility amount; *Log(Covenant)*, the natural logarithm of the covenant measure used in the analysis; *Log(Maturity)*, the natural logarithm of loan maturity in months; and *Collateral,* an indicator equal to one for secured loans. *Bank AI* equals one if the lead arranger filed at least one AI-related patent within five years preceding loan syndication. *Rel* equals one if the lead lender provided at least one loan to the borrower within the preceding five years. Panel A reports the direct effect of AI adoption. Panel B reports the interaction with relationship lending. The net effect for relationship borrowers equals the sum of the coefficients on *Bank AI* and *Bank AI* × *Rel*, and inference for this linear combination is obtained from an *F*-test. Economic magnitudes are computed as percentage changes for log variables and percentage point changes for indicator variables. All specifications include bank controls and loan controls, and borrower controls in the All and Public samples. Bank characteristics are bank size, liquidity, capital adequacy, profitability. Borrower characteristics are borrower size, book to market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, covenant. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates.***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

Panel A: DIRECT EFFECT OF AI ON NON-PRICE TERMS

	(1)	(2)	(3)	(4)
Dependent Variable:	Log(Loan Size)	Collateral	Log(Covenant)	Log(Maturity)
Sample	**All**	**All**	**All**	**All**
Bank AI	-0.035*	-0.006	-0.019**	-0.049***
	(0.019)	(0.009)	(0.008)	(0.015)
Implied econ. magnitude	-3.44%	-0.6 pp	-1.88%	-4.78%
Observations	64,672	64,693	63,809	60,340
R-squared	0.861	0.678	0.668	0.587
Sample	**Public**	**Public**	**Public**	**Public**
Bank AI	-0.046*	0.004	-0.027**	-0.041*
	(0.028)	(0.014)	(0.013)	(0.022)
Implied econ. magnitude	-4.50%	0.4 pp	-2.66%	-4.02%
Observations	25,888	25,892	25,892	24,297
R-squared	0.845	0.612	0.518	0.530
Sample	**Private**	**Private**	**Private**	**Private**
Bank AI	-0.019	-0.008	-0.03***	-0.058***
	(0.024)	(0.012)	(0.010)	(0.021)
Implied econ. magnitude	-1.88%	-0.8 pp	-2.96%	-5.64%
Observations	38,784	38,801	38,801	36,043
R-squared	0.879	0.729	0.678	0.635
Bank, Loan Controls	Yes	Yes	Yes	Yes
Borrower Controls	Yes (All| Public)	Yes (All| Public)	Yes (All| Public)	Yes (All| Public)
FEs: Borrower, Lender, Year	Yes	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes	Yes

Panel B: INTERACTION WITH RELATIONSHIP LENDING

	(1)	(2)	(3)	(4)
Dependent Variable:	Log(Loan Size)	Collateral	Log(Covenant)	Log(Maturity)
Sample	**All**	**All**	**All**	**All**
Bank AI	0.017	0.001	-0.025**	-0.031**
	(0.020)	(0.009)	(0.010)	(0.016)
Rel	0.135***	0.009	0.002	0.000
	(0.017)	(0.008)	(0.008)	(0.014)
Bank AI × Rel	-0.149***	-0.023***	-0.017*	-0.052***
	(0.018)	(0.009)	(0.009)	(0.015)
Net effect (relationship borrowers):				
Bank AI + Bank AI × Rel	-0.132***	-0.022*	-0.042***	-0.083***
[*t*-stat]	[-4.91]	[-1.73]	[-3.12]	[-3.78]
Implied econ. magnitude	-12.37%	-2.2 pp	-4.11%	-7.96%
Observations	64,672	64,693	64,693	60,340
R-squared	0.861	0.678	0.634	0.588
Sample	**Public**	**Public**	**Public**	**Public**
Bank AI	0.015	0.015	-0.014	-0.010
	(0.030)	(0.015)	(0.017)	(0.023)
Rel	0.194***	0.017	0.009	0.043**
	(0.025)	(0.012)	(0.014)	(0.019)
Bank AI × Rel	-0.178***	-0.023*	-0.035**	-0.080***
	(0.028)	(0.013)	(0.015)	(0.021)
Net effect (relationship borrowers):				
Bank AI + Bank AI × Rel	-0.163***	-0.008	-0.049**	-0.090***
[*t*-stat]	[-3.970]	[-0.400]	[-2.160]	[-2.890]
Implied econ. magnitude	-15.04%	-0.8 pp	-4.78%	-8.61%
Observations	25,888	25,892	25,400	24,297
R-squared	0.851	0.613	0.602	0.531
Sample	**Private**	**Private**	**Private**	**Private**
Bank AI	0.018	0.000	-0.03***	-0.054**
	(0.026)	(0.012)	(0.011)	(0.022)
Rel	0.084***	0.005	-0.013	-.049**
	(0.023)	(0.011)	(0.010)	(0.020)
Bank AI × Rel	-0.124***	-0.027**	0.000	-0.010
	(0.025)	(0.012)	(0.010)	(0.022)
Net effect (relationship borrowers):				
Bank AI + Bank AI × Rel	-0.106***	-0.027	-0.030**	-0.064**
[*t*-stat]	[-2.94]	[-1.59]	[-2.02]	[-2.06]
Implied econ. magnitude	-10.06%	-2.7 pp	-2.96%	-6.20%
Observations	38,784	38,801	38,801	36,043
R-squared	0.879	0.729	0.679	0.636
Bank, Loan Controls	Yes	Yes	Yes	Yes
Borrower Controls	Yes (All| Public)	Yes (All| Public)	Yes (All| Public)	Yes (All| Public)
FEs: Borrower, Lender, Year	Yes	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes	Yes

TABLE 10
Propensity Score Matching

This table reports regression estimates using propensity-score-matched (PSM) samples. We estimate a logit model for the probability that a loan facility is originated by an AI-adopting bank using bank size, profitability (ROA), and capital adequacy. Each treated facility (originated by an AI-adopting bank) is matched to its nearest neighbor in the control group using one-to-one nearest-neighbor matching with replacement based on the estimated propensity score. The dependent variable is *Loan Spread*, the all-in-drawn spread over LIBOR in basis points. *Bank AI* equals one if the lead arranger filed at least one AI-related patent within five years preceding loan syndication. All specifications include bank controls and loan controls, and borrower controls in the All and Public samples. Bank characteristics are bank size, liquidity, capital adequacy, profitability. Borrower characteristics are borrower size, book to market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, covenant. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. The definitions for all variables are in Table 1.Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. Significance at the 10 percent, 5 percent, and 1 percent levels is indicated by *, **, and ***, respectively.

	(1)	(2)	(3)
Sample	All	Public	Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread
Bank AI	11.647***	10.626***	8.923*
	(3.082)	(4.111)	(4.681)
Observations	21,944	10,090	11,854
R-squared	0.818	0.791	0.835
Bank, Loan Controls	Yes	Yes	Yes
Borrower Controls	Yes	No	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes

TABLE 11
Instrumental Variable (IV) Estimates

This table reports two-stage least squares (2SLS) estimates of the causal effect of bank AI adoption on loan spreads. The instrument is a Bartik-style interaction between banks' predetermined exposure to AI-intensive industries (measured over 1989–1994, before the sample period) and lagged growth in nonfinancial AI patents. $TechShock = Exposure_i \times Shock_{t-1}$, where $Exposure_i$ is the bank's share of lending to SIC codes 7371–7373 during 1989–1994, and $Shock_{t-1}$ is the lagged log-growth in AI patents granted to nonfinancial firms. Column (1) uses an indicator for positive *TechShock*; column (2) uses the continuous measure. Panel A reports IV first-stage estimates at the bank-month level; coefficients are marginal effects. Panel B reports second-stage estimates with Loan Spread (all-in-drawn spread over LIBOR, in basis points) as the dependent variable. All specifications include bank controls and loan controls, and borrower controls in the All and Public samples. Bank characteristics are bank size, liquidity, capital adequacy, profitability. Borrower characteristics are borrower size, book to market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, covenant. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. The definitions for all variables are in Table 1. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. Significance at the 10 percent, 5 percent, and 1 percent levels is indicated by *, **, and ***, respectively.

Panel A: IV FIRST STAGE

	(1)	(2)
Sample	All	All
Dependent Variable:	Bank AI	Bank AI
TechShock (Indicator)	0.636***	
	(0.049)	
TechShock (Continuous)		3.572**
		(1.821)
Observations	11,144	11,144
R-squared	0.307	0.294
Bank Controls	Yes	Yes
Year FE	Yes	Yes

Panel B: IV LAST STAGE

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	All	All	Public	Public	Private	Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread	Loan Spread
Instrument:	Indicator	Continuous	Indicator	Continuous	Indicator	Continuous
Bank AI (Instrumented)	80.299***	107.893***	90.450***	115.688***	69.528	92.788***
	(27.974)	(17.722)	(29.997)	(20.456)	(61.439)	(35.107)
Observations	51,630	51,630	21,832	21,832	29,798	29,798
R-squared	0.077	0.068	0.154	0.143	0.065	0.059
Bank, Loan Controls	Yes	Yes	Yes	Yes	Yes	Yes
Borrower Controls	Yes	Yes	Yes	Yes	—	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes	Yes	Yes	Yes
KP Weak Identification	555.3***	1428***	387.7***	866***	146***	456***
KP UnderIdentification	700.8***	1764***	461.8***	1005***	193.4***	596***

TABLE 12
Robustness: AI vs. Non-AI Parents (Falsification)

This table reports a falsification test comparing the effects of AI patents versus non-AI patents on loan spreads. *Bank AI Patents* equals one if the lead arranger filed at least one patent with AI content score ≥ 50% (per the USPTO AIPD classification) within five years preceding loan syndication. *Bank Non-AI Patents* equals one if the lead arranger filed at least one patent with AI content score < 10% within five years preceding loan syndication. The dependent variable is the all-in-drawn spread over LIBOR in basis points. If our results reflect general bank technological sophistication rather than AI-specific information advantages, both measures should have similar effects. The difference row reports the *AI* coefficient minus the *Non-AI* coefficient with t-values. For public firms, AI patents have significantly larger effects than non-AI patents, supporting AI-specific mechanisms. All specifications include bank controls and loan controls, and borrower controls in the All and Public samples. Bank characteristics are bank size, liquidity, capital adequacy, profitability. Borrower characteristics are borrower size, book to market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, covenant. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. The definitions for all variables are in Table 1. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. Significance at the 10 percent, 5 percent, and 1 percent levels is indicated by *, **, and ***, respectively.

Sample	(1) All	(2) Public	(3) Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread
Bank AI Patents	9.250*	11.37*	7.714
	(4.765)	(5.867)	(6.413)
Bank Non-AI Patents	3.770	3.428	6.077
	(3.003)	(2.948)	(6.522)
Difference (Bank AI − Bank Non-AI):	5.48	7.94	1.63
[*t*-stat]	[0.76]	[1.16]	[0.13]
Observations	52,041	22,145	29,896
R-squared	0.678	0.661	0.684
Bank, Loan Controls	Yes	Yes	Yes
Borrower Controls	Yes	Yes	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes

TABLE 13
Additional Robustness Tests

This table reports additional robustness tests. Panel B excludes the three largest lenders (by total assets) to ensure results are not driven by a small number of dominant banks; effects are larger when excluding these institutions. Panel C uses AI intensity (share of AI patents out of total patents filed by the bank in the preceding five years) rather than an indicator, showing that more intensive AI adoption is associated with higher spreads. Panel C reports baseline results with standard errors clustered at the bank level. The dependent variable is *Loan Spread*, the all-in-drawn spread over LIBOR in basis points. All specifications include bank controls and loan controls, and borrower controls in the All and Public samples. Bank characteristics are bank size, liquidity, capital adequacy, profitability. Borrower characteristics are borrower size, book to market, leverage, ROA, and tangibility. Loan characteristics are loan size, maturity, collateral, covenant. All specifications include borrower, lender, and year fixed effects, as well as loan type and loan purpose fixed effects. The definitions for all variables are in Table 1. Heteroskedasticity-robust standard errors are reported in parentheses beneath the coefficient estimates. Significance at the 10 percent, 5 percent, and 1 percent levels is indicated by *, **, and ***, respectively.

Panel A: EXCLUDE TOP3 LENDERS

	(1)	(2)	(3)
Sample	All	Public	Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread
Bank AI	9.844***	11.960**	7.163
	(4.287)	(5.957)	(6.486)
Observations	12,078	4,718	7,359
R-squared	0.767	0.733	0.779
Bank, Loan Controls	Yes	Yes	Yes
Borrower Controls	Yes	No	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes

Panel B: AI INTENSITY (SHARE OF AI PATENTS)

	(1)	(2)	(3)
Sample	All	Public	Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread
Bank AI	6.088**	7.879**	4.945
	(2.833)	(3.849)	(4.303)
Observations	52,041	22,145	29,896
R-squared	0.678	0.661	0.684
Bank, Loan Controls	Yes	Yes	Yes
Borrower Controls	Yes	No	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes

Panel C: BANK-CLUSTERED STANDARD ERRORS

	(1)	(2)	(3)
Sample	All	Public	Private
Dependent Variable:	Loan Spread	Loan Spread	Loan Spread
Bank AI	10.47**	12.58***	9.329*
	(4.137)	(5.877)	(5.031)
Observations	52,041	22,145	29,896
R-squared	0.678	0.661	0.684
Bank, Loan Controls	Yes	Yes	Yes
Borrower Controls	Yes	No	—
FEs: Borrower, Lender, Year	Yes	Yes	Yes
FEs: Loan Type, Loan Purpose	Yes	Yes	Yes

ONLINE APPENDIX (OA): SUPPLEMENTARY MATERIALS

FIGURE OA.1:
Evolution of Bank AI Patenting over Time

This figure plots the annual count of AI-related patents filed by banks by application year using the USPTO AI patents dataset. Patents are attributed to banks based on the assignee. The blue series reports annual counts, and the red series reports the three-year moving average. Counts in the most recent years may be mechanically lower due to right-censoring and processing lags in patent publication and grant timing. The *x*-axis reports application year and the *y*-axis reports the annual number of bank AI patents, a proxy for the intensity of bank AI innovation over the sample period.

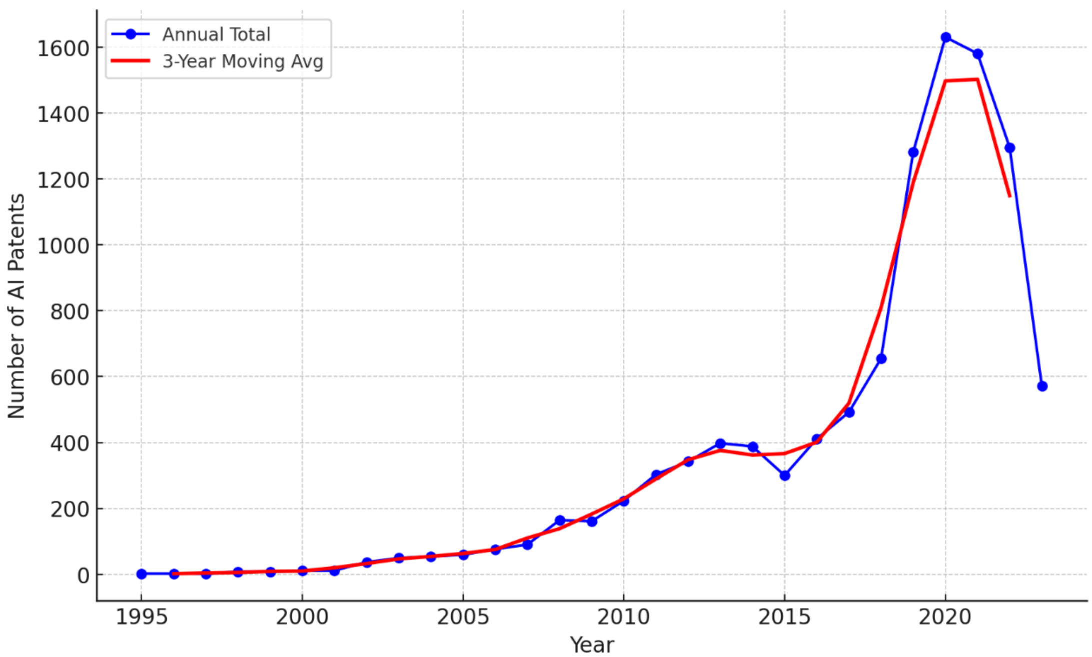

FIGURE OA.2:
Dynamic Effects of Bank AI Adoption on Loan Spreads:
Sun-Abraham (2021) Event Study

This figure plots interaction-weighted (IW) coefficient estimates from a Sun and Abraham (2021) event study of the dynamic effects of bank AI adoption on loan spreads. The event is defined as the first year a bank's *AI_past_5yrs_dummy* equals one. The sample is restricted to switchers and never-treated banks; always-treated banks are excluded from the control group. The omitted reference period is t =−1. Each coefficient measures the average deviation in the all-in-drawn spread (in basis points) relative to the year immediately preceding first AI adoption, with 95% confidence intervals constructed using standard errors clustered at the lender level. Fixed effects include borrower and lender identifiers. Year fixed effects are omitted due to near-collinearity with event-time indicators arising from the temporal clustering of AI adoption; the pre-treatment pattern therefore reflects aggregate trends in market-wide spread levels rather than anticipation effects. See Section OA.1 for discussion. The x-axis reports years relative to first AI adoption and the y-axis reports the IW-weighted coefficient estimate in basis points.

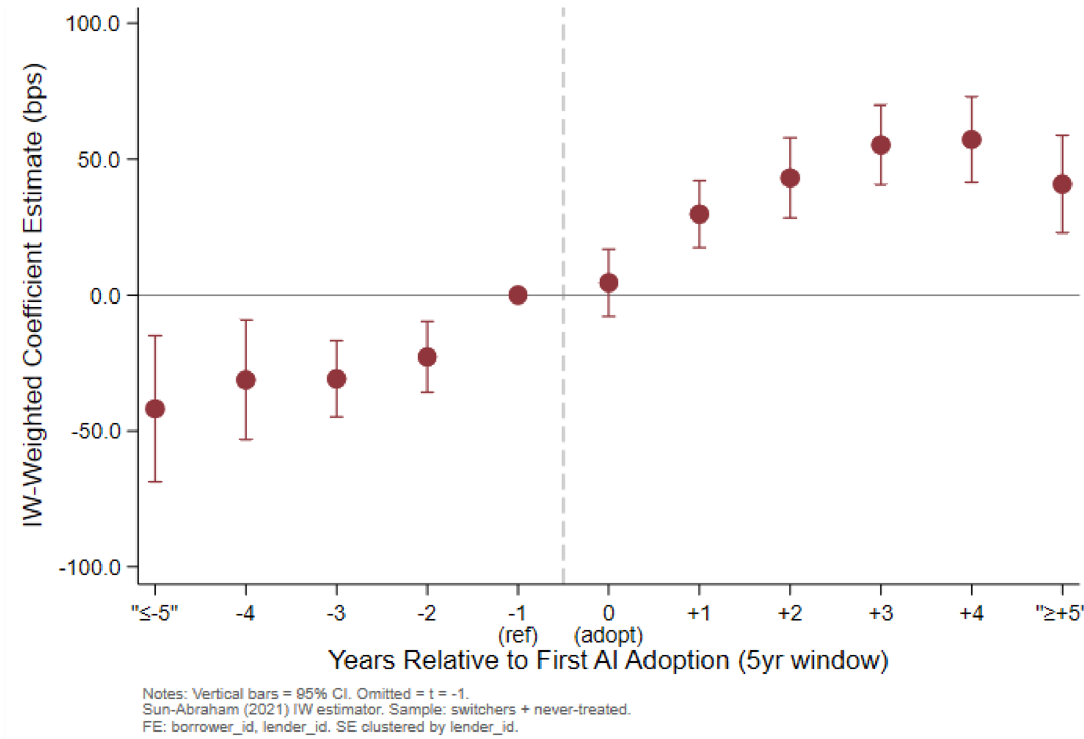

OA.1 Event Study Evidence on Dynamic Effects

To complement the cross-sectional identification in our main specifications and the Bartik-style IV, we examine the dynamic response of loan spreads around banks' first AI adoption using an event study design. We define the event as the first year a bank's *AI_past_5yrs_dummy* equals one and estimate dynamic treatment effects using the Sun and Abraham (2021) interaction-weighted (IW) estimator, which is robust to heterogeneous treatment effects across cohorts and avoids the negative-weighting problem that affects standard two-way fixed effects event studies (Callaway and Sant'Anna (2021); de Chaisemartin and D'Haultfœuille (2020)). The sample is restricted to switchers, banks that adopt AI at some point during our sample period and never-treated banks, excluding always-treated banks from the control group. Standard errors are clustered at the lender level and fixed effects include borrower and lender identifiers. We omit year fixed effects in this specification; as we discuss below, year fixed effects absorb the primary source of identifying variation because AI adoption is substantially clustered in calendar time, and their inclusion mechanically attenuates all event-time coefficients toward zero.

Figure OA.2 reports the IW-weighted coefficient estimates for event times $t \leq -5$ through $t \geq +5$, with $t = -1$ as the omitted reference period. The post-treatment pattern is clear and economically striking. Loan spreads rise modestly and insignificantly at $t = 0$ (4.5 basis points, the year of first adoption), then increase sharply and monotonically: 29.8 basis points at $t = +1$, 43.1 basis points at $t = +2$, 55.2 basis points at $t = +3$, and 57.3 basis points at $t = +4$, all significant at the 1% level. The slight decline to 40.8 basis points in the ≥+5 bin likely reflects endpoint compression and the pooling of heterogeneous long-run effects rather than a genuine mean reversion. The gradual buildup from $t = 0$ to $t = +3/+4$ is consistent with the informational hold-up mechanism: proprietary borrower intelligence accumulates over successive lending cycles, and the bank's ability to extract rents increases as its information advantage deepens relative to outside lenders. The dynamic pattern also rules out a mechanical explanation based on contemporaneous market conditions, since the spread increase builds steadily over multiple years following adoption.

We acknowledge an important limitation of this design. The pre-treatment coefficients are negative and statistically significant: −22.8 basis points at $t = -2$, −30.8 basis points at $t = -3$, and −31.2 basis points at $t = -4$, violating the strict parallel trends assumption required for a causal interpretation. We attribute this pattern to the absence of year fixed effects, which cannot be included in this specification because AI adoption is heavily concentrated in calendar time, while some banks filed AI patents as early as the mid-1990s, the large majority of adoption events and virtually all of the intensive-margin variation occur in the post-2015 surge period, creating near-collinearity between year dummies and event-time indicators in the full sample. When year fixed effects are added, coefficients, both pre- and post-treatment, lose significance almost entirely, confirming that year fixed effects and event-time dummies span the same variation in this setting. The significant pre-period coefficients therefore reflect aggregate trends in market-wide spread levels over our 1995–2023 sample period rather than anticipation effects or changes in bank behavior prior to adoption.

Consistent with this interpretation, the pre-trend pattern is monotonically negative and increasing in magnitude as we move further from $t = -1$ the further back in calendar time, the lower were aggregate spreads relative to the omitted period, which is precisely the signature of a secular upward trend in market-wide spreads being absorbed into the pre-period coefficients rather than reflecting bank-specific behavior prior to adoption. For these reasons, we treat the event study as supplementary dynamic evidence rather than a primary identification test. Our Bartik-style IV strategy, which exploits variation in predetermined industry exposure interacted with aggregate AI innovation shocks, is not subject to these concerns and provides the cleanest causal estimates in the paper.